

Extreme NetIron Multicast Configuration Guide, 06.3.00

Supporting NetIron OS 06.3.00

© 2018, Extreme Networks, Inc. All Rights Reserved.

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 are the property of their respective owners. For additional information on Extreme Networks Trademarks please see www.extremenetworks.com/company/legal/trademarks. Specifications and product availability are subject to change without notice.

Contents

Preface	7
Document conventions.....	7
Notes, cautions, and warnings.....	7
Text formatting conventions.....	7
Command syntax conventions.....	8
Extreme resources.....	8
Document feedback.....	8
Contacting Extreme Technical Support.....	9
About This Document	11
What's new in this document.....	11
Supported hardware and software.....	11
Supported software.....	11
How command information is presented in this guide.....	12
IPv4 Multicast VLAN Traffic Reduction	13
IP multicast traffic reduction.....	13
Configuration requirements.....	13
Configuring IP multicast traffic reduction.....	13
PIM SM traffic snooping.....	16
Multicast traffic reduction per VLAN or VPLS instance.....	20
IP Multicast CPU Protection.....	23
Static IGMP membership.....	26
Displaying IP multicast information.....	28
IPv6 Multicast VLAN Traffic Reduction	31
IPv6 Multicast Listener Discovery snooping.....	31
Configuring IPv6 multicast routing or snooping.....	31
Enabling IPv6 multicast traffic reduction.....	31
PIM-SM traffic snooping.....	33
Configuring IPv6 MLD snooping on a per-VLAN basis.....	37
Displaying IPv6 multicast information.....	40
IPv4 Multicast Routing	43
Overview of IP multicasting.....	43
Multicast terms.....	43
Changing global IP multicast parameters.....	44
Concurrent support for multicast routing and snooping.....	44
Defining the maximum number of PIM cache entries.....	44
Defining the maximum number of multicast VRF CAM entries.....	45
Defining the maximum number of IGMP group addresses.....	45
Changing IGMP V1 and V2 parameters.....	46
Mtrace overview.....	47
Mtrace components.....	48
Configuring mtrace.....	49
Support for Multicast Multi-VRF.....	49
System max parameter changes.....	49
Show and clear command support.....	50

Adding an interface to a multicast group.....	50
Multicast non-stop routing.....	51
Configuration considerations.....	51
Configuring multicast non-stop routing.....	51
Displaying the multicast NSR status.....	52
Passive Multicast Route Insertion (PMRI)	53
Configuring PMRI.....	53
Displaying hardware-drop.....	53
IP multicast boundaries.....	54
Configuration considerations.....	54
Configuring multicast boundaries.....	54
Displaying multicast boundaries.....	55
Performing IPv4 Multicast RPF shortcut using LSP paths.....	56
Multicast ECMP support.....	58
Limitations and prerequisites.....	60
Enabling multicast fast convergence.....	61
Configuring Layer 3 Multicast filter for the hardware.....	62
Limitations and pre-requisites.....	62
Configuring the Layer 3 Multicast filter.....	63
Displaying Multicast filter for the hardware.....	64
PIM Dense	65
Initiating PIM multicasts on a network.....	65
Pruning a multicast tree.....	65
Grafts to a multicast tree.....	67
PIM DM versions.....	68
Configuring PIM DM	68
Failover time in a multi-path topology.....	72
Modifying the TTL threshold.....	72
Configuring a DR priority.....	72
Displaying basic PIM Dense configuration information.....	73
Displaying all multicast cache entries in a pruned state.....	74
Displaying all multicast cache entries.....	74
Multicast PIM neighbor filter.....	78
PIM Sparse	80
PIM Sparse device types.....	81
RP paths and SPT paths.....	82
Configuring PIM Sparse.....	82
ACL based RP assignment.....	86
Route selection precedence for multicast.....	87
PIM multinet.....	90
Multicast Outgoing Interface (OIF) list optimization.....	91
Displaying PIM Sparse configuration information and statistics.....	91
Clearing the PIM forwarding cache.....	105
Displaying PIM traffic statistics.....	105
Clearing the PIM message counters.....	106
Displaying PIM counters.....	106
Configuring Multicast Source Discovery Protocol (MSDP).....	107
Peer Reverse Path Forwarding (RPF) flooding.....	109
Source Active caching.....	109
Configuring MSDP.....	109

Disabling an MSDP peer.....	111
Designating the interface IP address as the RP IP address.....	111
Filtering MSDP source-group pairs.....	111
Filtering incoming and outgoing Source-Active messages.....	112
Filtering advertised Source-Active messages.....	113
Displaying MSDP information.....	114
Displaying MSDP RPF-Peer.....	118
Displaying MSDP Peer.....	119
Displaying MSDP VRF RPF-Peer.....	119
Clearing MSDP information.....	119
Configuring MSDP mesh groups	120
Configuring MSDP mesh group.....	122
MSDP Anycast RP.....	122
Configuring MSDP Anycast RP.....	123
Example.....	123
PIM Anycast RP.....	126
Configuring PIM Anycast RP.....	126
PIM over MCT intermediate router functionality.....	128
MCT peer as intermediate Upstream router	129
MCT peer as intermediate Downstream router	131
MCT peers as PIM Anycast RP.....	132
Source directly connected to CEP on MCT VLAN.....	133
Multi tier MCT.....	134
Limitations.....	136
Enabling PIM over MCT scaling optimization.....	136
Displaying IGMP and MLD cluster group information.....	137
Displaying MCT PIM Counters.....	137
Configuring a static multicast route.....	139
Configuring a static multicast route within a VRF.....	140
IGMP V3.....	141
Default IGMP version.....	142
Compatibility with IGMP V1 and V2.....	142
Globally enabling the IGMP version	142
Enabling the IGMP version per interface setting	142
Enabling the IGMP version on a physical port within a virtual routing interface	143
Enabling membership tracking and fast leave.....	143
Creating a static IGMP group.....	144
Setting the query interval.....	144
Setting the group membership time.....	144
Setting the maximum response time.....	145
Displaying IGMPv3 information.....	145
Displaying IGMP group status.....	145
Clearing the IGMP group membership table	147
Displaying static IGMP groups.....	148
Clearing IGMP traffic statistics	150
Source-specific multicast.....	151
Configuring PIM SSM group range.....	151
Configuring multiple SSM group ranges.....	151
IGMPv2 SSM mapping.....	153
IPv6 Multicast Routing.....	157

IPv6 PIM Sparse	157
PIM Sparse router types.....	158
RP paths and SPT paths.....	158
RFC 3513 and RFC 4007 compliance for IPv6 multicast scope-based forwarding.....	158
Configuring PIM Sparse.....	159
IPv6 PIM-Sparse mode.....	159
Configuring IPv6 PIM-SM on a virtual routing interface.....	159
Enabling IPv6 PIM-SM for a specified VRF.....	160
Configuring BSRs	160
Route selection precedence for multicast.....	164
Enabling Source-specific Multicast.....	167
Configuring a DR priority.....	168
Passive Multicast Route Insertion.....	169
Displaying PIM Sparse configuration information and statistics.....	170
Clearing the IPv6 PIM forwarding cache.....	181
Clearing the IPv6 PIM message counters.....	181
Updating PIM Sparse forwarding entries with a new RP configuration.....	182
Clearing the IPv6 PIM traffic	182
Setting the maximum number of IPv6 multicast routes supported	182
Defining the maximum number of IPv6 PIM cache entries.....	183
Defining the maximum number of IPv6 multicast VRF CAM entries for all VRFs.....	183
PIM Anycast RP.....	183
Configuring PIM Anycast RP.....	184
Multicast Listener Discovery and source-specific multicast protocols.....	186
Enabling MLDv2.....	186
Configuring MLD parameters for default and non-default VRFs.....	186
Configuring MLD parameters at the interface level.....	189
Displaying MLD information.....	190
Clearing IPv6 MLD traffic.....	194
Clearing the IPv6 MLD group membership table cache.....	194

Preface

- Document conventions..... 7
- Extreme resources..... 8
- Document feedback..... 8
- Contacting Extreme Technical Support..... 9

Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Extreme technical documentation.

Notes, cautions, and warnings

Notes, cautions, and warning statements may be used in this document. They are listed in the order of increasing severity of potential hazards.

NOTE

A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

ATTENTION

An Attention statement indicates a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.



CAUTION

A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.



DANGER

A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used to highlight specific words or phrases.

Format	Description
bold text	Identifies command names. Identifies keywords and operands. Identifies the names of GUI elements.
<i>italic text</i>	Identifies text to enter in the GUI. Identifies emphasis. Identifies variables.
Courier font	Identifies document titles. Identifies CLI output.

Format	Description
	Identifies command syntax examples.

Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
bold text	Identifies command names, keywords, and command options.
<i>italic text</i>	Identifies a variable.
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member[member...]</i> .
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Extreme resources

Visit the Extreme website to locate related documentation for your product and additional Extreme resources.

White papers, data sheets, and the most recent versions of Extreme software and hardware manuals are available at www.extremenetworks.com. Product documentation for all supported releases is available to registered users at www.extremenetworks.com/support/documentation.

Document feedback

Quality is our first concern at Extreme, and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you.

You can provide feedback in two ways:

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

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

Contacting Extreme Technical Support

As an Extreme customer, you can contact Extreme Technical Support using one of the following methods: 24x7 online or by telephone. OEM customers should contact their OEM/solution provider.

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

- [GTAC \(Global Technical Assistance Center\)](#) for immediate support
 - Phone: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact.
 - Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- [GTAC Knowledge](#) - Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- [The Hub](#) - A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- [Support Portal](#) - Manage cases, downloads, service contracts, product licensing, and training and certifications.

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

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

About This Document

- [What's new in this document.....](#) 11
- [Supported hardware and software.....](#) 11
- [How command information is presented in this guide.....](#) 12

What's new in this document

There has been no enhancement to this guide for the NI 6.3.00 software release.

NOTE

On October 30, 2017, Extreme Networks, Inc. acquired the SLX-OS product line from Brocade Communications Systems, Inc. This transitional release includes references to both companies.

Supported hardware and software

End of Support for ExtremeSwitching CES 2000 Series devices

Beginning with NetIron OS 06.3.00 and later, the ExtremeSwitching CES 2000 Series devices are not supported. Refer to the [End of Sale and End of Support](#) page for additional information.

The hardware platforms in the following table are supported by this release of this guide.

TABLE 1 Supported devices

ExtremeRouting XMR Series	ExtremeRouting MLX Series	ExtremeRouting CER 2000 Series
XMR 4000	MLX-4	CER 2024C
XMR 8000	MLX-8	CER-RT 2024C
XMR 16000	MLX-16	CER 2024F
XMR 32000	MLX-32	CER-RT 2024F
	MLXe-4	CER 2048C
	MLXe-8	CER-RT 2048C
	MLXe-16	CER 2048CX
	MLXe-32	CER-RT 2048CX
		CER 2048F
		CER-RT 2048F
		CER 2048FX
		CER-RT 2048FX

Supported software

For the complete list of supported features and the summary of enhancements and configuration notes for this release, refer to the *Extreme NetIron Release Notes*.

How command information is presented in this guide

Starting with Extreme NetIron 5.6.00, command syntax and parameter descriptions are removed from commands that are referenced in configuration tasks. To find the full description of a specific command, including all required and optional keywords and variables, refer to the *Extreme NetIron Command Reference* for your software release.

IPv4 Multicast VLAN Traffic Reduction

- [IP multicast traffic reduction.....](#) 13

IP multicast traffic reduction

In Layer 2 mode, by default, the Extreme device forwards all IP multicast traffic out all ports except the port on which the traffic was received. Forwarding decisions are based on the Layer 2 information in the packets. To reduce multicast traffic through the device, you can enable IP Multicast Traffic Reduction. When this feature is enabled, forwarding decisions are made in hardware, based on multicast group. The device will forward multicast traffic only on the ports attached to multicast group members, instead of forwarding all multicast traffic to all ports.

By default, the device broadcasts traffic addressed to an IP multicast group that does not have any entries in the IGMP table. When you enable IP Multicast Traffic Reduction, the device determines the ports that are attached to multicast group members based on entries in the IGMP table. The IGMP table entries are created when the VLAN receives a group membership report for a group. Each entry in the table consists of an IP multicast group address and the ports from which the device has received Group Membership reports.

When the device receives traffic for an IP multicast group, the device looks in the IGMP table for an entry corresponding to that group. If the device finds an entry, it forwards the group traffic out the ports listed in the corresponding entries, as long as the ports are members of the same VLAN. If the table does not contain an entry corresponding to the group, or if the port is a member of the default VLAN, the device broadcasts the traffic.

Configuration requirements

Consider the following configuration requirements and application notes:

- The IP Multicast Traffic Reduction feature is applicable to Layer 2 mode only.
- If the **route-only** feature is enabled on the Extreme device, then IP Multicast Traffic Reduction will not be supported.
- This feature is not supported on the default VLAN of the Extreme device.
- When one or more Extreme devices are running Layer 2 IP Multicast Traffic reduction, configure one of the devices for active IGMP and leave the other devices configured for passive IGMP. However, if the IP multicast domain contains a multicast-capable router, configure all the Extreme devices for passive IGMP and allow the router to actively send the IGMP queries.
- IP multicast traffic reduction and PIM SM Traffic Snooping are supported on the Extreme device.

Configuring IP multicast traffic reduction

When you enable IP Multicast Traffic Reduction, you also can configure the following features:

- **IGMP mode** - When you enable IP Multicast Traffic Reduction, the device passively listens for IGMP Group Membership reports by default. If the multicast domain does not have a router to send IGMP queries to elicit these Group Membership reports, you can enable the device to actively send the IGMP queries. The IGMP passive mode is also known as IGMP snooping and facilitates IP Multicast Traffic Reduction.
- **Query interval** - The query interval specifies how often the device sends Group Membership queries. This query interval applies only to the active IGMP mode. The default is 60 seconds. You can change the interval to a value from 10 - 600 seconds.
- **Age interval** - The age interval specifies how long an IGMP group can remain in the IGMP group table without the device receiving a Group Membership report for the group. If the age interval expires before the device receives another Group

Membership report for the group, the device removes the entry from the table. The default is 140 seconds. You can change the interval to a value from 10 - 1220 seconds.

Furthermore, when you enable IP Multicast Traffic Reduction, the device forwards all IP multicast traffic by default, but you can enable the device to do the following:

- Forward IP multicast traffic only for groups for which the device has received a Group Membership report.
- Drop traffic for all other groups.

The following sections describe how to configure IP multicast traffic reduction and PIM SM Traffic Snooping parameters on a Extreme device.

Enabling IP multicast traffic reduction

To enable IP Multicast Traffic Reduction, enter the following command.

```
device(config)# ip multicast
```

Syntax: [no] ip multicast active | passive

When you enable IP multicast on a Extreme device, all ports on the device are configured for IGMP.

The **active** mode enables all ports to send IGMP queries and receive IGMP reports. I

The **passive** mode enables all ports to receive IGMP queries.

IP Multicast Traffic Reduction cannot be disabled on individual ports of a Extreme device. IP Multicast Traffic Reduction must be disabled globally by entering the **no ip multicast** command.

To verify that IP Multicast Traffic Reduction is enabled, enter the following command at any level of the CLI.

```
device(config)# show ip multicast
IP multicast is enabled - Active
```

Syntax: show ip multicast

Changing the IGMP mode

When you enable IP Multicast Traffic Reduction on the device, IGMP also is enabled. The device uses IGMP to maintain a table of the Group Membership reports received by the device. You can use active or passive IGMP mode. There is no default mode.

The active and passive IGMP modes are described as follows:

- **Active** - When active IGMP mode is enabled, a Extreme device actively sends out IGMP queries to identify IP multicast groups on the network and makes entries in the IGMP table based on the Group Membership reports received from the network.

NOTE

Routers in the network generally handle this operation. Use the active IGMP mode only when the device is in a stand-alone Layer 2 Switched network with no external IP multicast router attachments. In this case, enable the active IGMP mode on only one of the devices and leave the other devices configured for passive IGMP mode.

- **Passive** - When passive IGMP mode is enabled, the device listens for IGMP Group Membership reports but does not send IGMP queries. The passive mode is sometimes called "IGMP snooping". Use this mode when another device in the network is actively sending queries.

To enable active IGMP, enter the following command.

```
device(config)# ip multicast active
```

Syntax: [no] ip multicast active | passive

To enable passive IGMP, enter the following command.

```
device(config)# ip multicast passive
```

Modifying the query interval

If IP Multicast Traffic Reduction is set to active mode, you can modify the query interval, which specifies how often a Extreme device enabled for active IP Multicast Traffic Reduction sends group membership queries.

NOTE

The query interval applies only to the active mode of IP Multicast Traffic reduction.

To modify the query interval, enter a command such as the following.

```
device(config)# ip multicast query-interval 120
```

Syntax: `[no] ip multicast query-interval interval`

The interval parameter specifies the interval between queries. You can specify a value from 10 - 600 seconds. The default is 125 seconds.

Modifying the age interval

When the device receives a Group Membership report, the device makes an entry in the IGMP group table for the group in the report. The age interval specifies how long the entry can remain in the table without the device receiving another Group Membership report.

To modify the age interval, enter a command such as the following.

```
device(config)# ip multicast age-interval 280
```

Syntax: `[no] ip multicast age-interval interval`

The interval parameter specifies the interval between queries. You can specify a value from 30 - 1220 seconds. The default is 140 seconds.

Filtering multicast groups

By default, the Extreme device forwards multicast traffic for all valid multicast groups. You can configure a Extreme device to filter out all multicast traffic for groups other than the ones for which the device has received Group Membership reports.

When the device starts up, it forwards all multicast groups even though multicast traffic filters are configured. This process continues until the device receives a group membership report. Once the group membership report is received, the device drops all multicast packets for groups other than the ones for which the device has received the group membership report.

To enable IP multicast filtering, enter the following command.

```
device(config)# ip multicast filter
```

Syntax: `[no] ip multicast filter`

NOTE

When IGMP snooping is enabled on the multicast instance, the traffic will be forwarded to router ports although multicast filter is configured. Also, the ip multicast filter command is used only to avoid flooding. In case of PIM snooping, traffic is dropped since there are no router ports.

PIM SM traffic snooping

By default, when a Extreme device receives an IP multicast packet, the device does not examine the multicast information in the packet. Instead, the device simply forwards the packet out all ports except the port that received the packet. In some networks, this method can cause unnecessary traffic overhead in the network. For example, if the Extreme device is attached to only one group source and two group receivers, but has devices attached to every port, the device forwards group traffic out all ports in the same broadcast domain except the port attached to the source, even though there are only two receivers for the group.

PIM SM traffic snooping eliminates the superfluous traffic by configuring the device to forward IP multicast group traffic only on the ports that are attached to receivers for the group.

PIM SM traffic snooping requires IP multicast traffic reduction to be enabled on the device. IP multicast traffic reduction configures the device to listen for IGMP messages. PIM SM traffic snooping provides a finer level of multicast traffic control by configuring the device to listen specifically for PIM SM join and prune messages sent from one PIM SM router to another through the device.


NOTE

This feature applies only to PIM SM version 2 (PIM V2).

Application examples

Figure 1 shows an example application of the PIM SM traffic snooping feature. In this example, a device is connected through an IP router to a PIM SM group source that is sending traffic for two PIM SM groups. The device also is connected to a receiver for each of the groups.

FIGURE 1 PIM SM traffic reduction in enterprise network



When PIM SM traffic snooping is enabled, the device starts listening for PIM SM join and prune messages and IGMP group membership reports. Until the device receives a PIM SM join message or an IGMP group membership report, the device forwards IP multicast traffic out all ports. Once the device receives a join message or group membership report for a group, the device forwards subsequent traffic for that group only on the ports from which the join messages or IGMP reports were received.

In this example, the router connected to the receiver for group 239.255.162.1 sends a join message toward the group's source. Since PIM SM traffic snooping is enabled on the device, the device examines the join message to learn the group ID, then makes a forwarding entry for the group ID and the port connected to the receiver's router. The next time the device receives traffic for 239.255.162.1 from the group's source, the device forwards the traffic only on port 5/1, since that is the only port connected to a receiver for the group.

Notice that the receiver for group 239.255.162.69 is directly connected to the device. As result, the device does not see a join message on behalf of the client. However, since IP multicast traffic reduction also is enabled, the device uses the IGMP group membership report from the client to select the port for forwarding traffic to group 239.255.162.69 receivers.

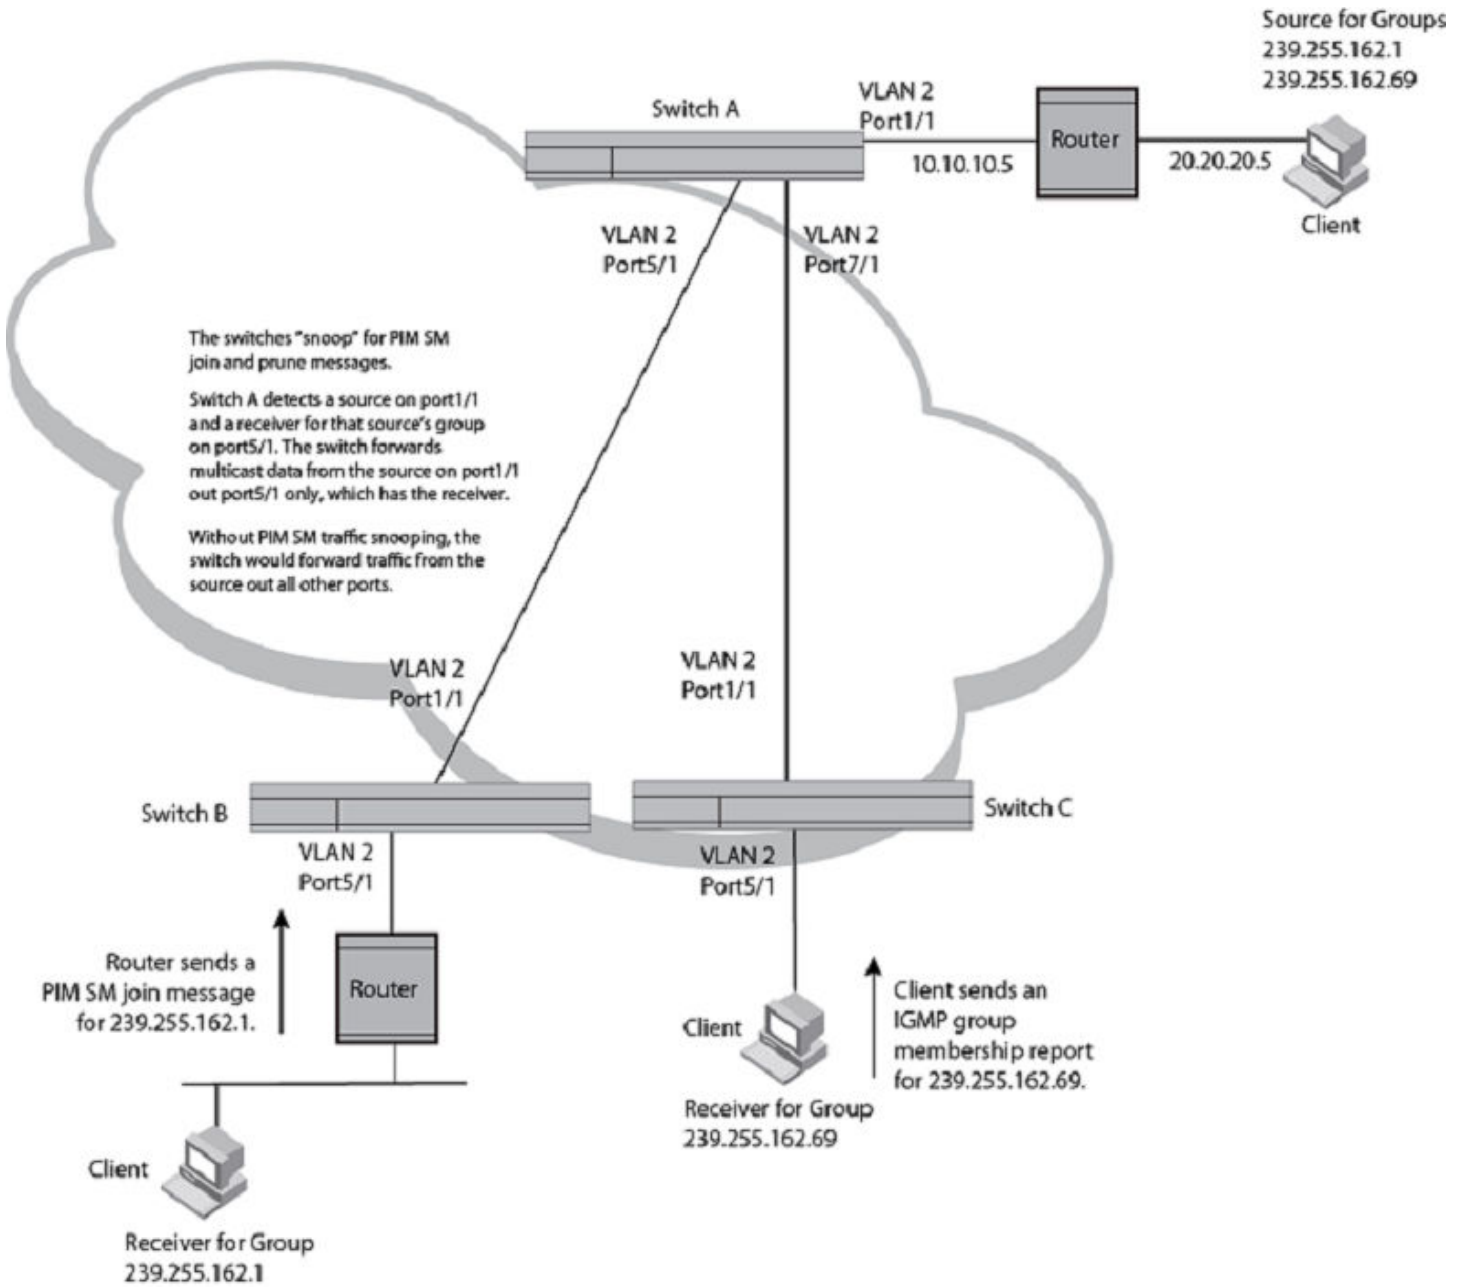
The IP multicast traffic reduction feature and the PIM SM traffic snooping feature together build a list of groups and forwarding ports for the VLAN. The list includes PIM SM groups learned through join messages as well as MAC addresses learned through IGMP group membership reports. In this case, even though the device never sees a join message for the receiver for group 239.255.162.69, the device nonetheless learns about the receiver and forwards group traffic to the receiver.

The device stops forwarding IP multicast traffic on a port for a group if the port receives a prune message for the group.

Notice that the ports connected to the source and the receivers are all in the same port-based VLAN on the device. This is required for the PIM SM snooping feature. The feature also requires the source and the downstream router to be on different IP subnets, as shown in [Figure 1](#).

[Figure 2](#) shows another example application for PIM SM traffic snooping. This example shows devices on the edge of a Global Ethernet cloud (a Layer 2 Packet over SONET cloud). Assume that each device is attached to numerous other devices such as other Extreme devices.

FIGURE 2 PIM SM traffic reduction in Global Ethernet environment



The devices on the edge of the Global Ethernet cloud are configured for IP multicast traffic reduction and PIM SM traffic snooping. Although this application uses multiple devices, the feature has the same requirements and works the same way as it does on a single device.

Configuration requirements

Consider the following configuration requirements:

- IP multicast traffic reduction must be enabled on the device that will be running PIM SM snooping. The PIM SM traffic snooping feature requires IP multicast traffic reduction.

NOTE

Use the passive mode of IP multicast traffic reduction instead of the active mode. The passive mode assumes that a router is sending group membership queries as well as join and prune messages on behalf of receivers. The active mode configures the device to send group membership queries.

- All the device ports connected to the source and receivers or routers must be in the same port-based VLAN.
- The PIM SM snooping feature assumes that the group source and the device are in different subnets and communicate through a router. The source must be in a different IP subnet than the receivers. A PIM SM router sends PIM join and prune messages on behalf of a multicast group receiver only when the router and the source are in different subnets. When the receiver and source are in the same subnet, they do not need the router in order to find one another. They find one another directly within the subnet. The device forwards all IP multicast traffic by default. Once you enable IP multicast traffic reduction and PIM SM traffic snooping, the device initially blocks all PIM SM traffic instead of forwarding it. The device forwards PIM SM traffic to a receiver only when the device receives a join message from the receiver. Consequently, if the source and the downstream router are in the same subnet, and PIM SM traffic snooping is enabled, the device blocks the PIM SM traffic and never starts forwarding the traffic. This is because the device never receives a join message from the downstream router for the group. The downstream router and group find each other without a join message because they are in the same subnet.

NOTE

If the "route-only" feature is enabled on a Extreme device, PIM SM traffic snooping will not be supported.

Enabling PIM SM traffic snooping

To enable PIM SM traffic snooping, enter the following commands at the global CONFIG level of the CLI.

```
device(config)# ip multicast
device(config)# ip pimsm-snooping
```

The first command enables IP multicast traffic reduction. This feature is similar to PIM SM traffic snooping but listens only for IGMP information, not PIM SM information. You must enable both IP multicast traffic reduction and PIM SM traffic snooping to enable the device to listen for PIM SM join and prune messages.

Syntax: `[no] ip multicast [active | passive]`

This command enables IP multicast traffic reduction. The **active** | **passive** parameter specifies the mode. The PIM SM traffic snooping feature assumes that the network has routers that are running PIM SM.

Syntax: `[no] ip pimsm-snooping`

This command enables PIM SM traffic snooping.

To disable the feature, enter the following command.

```
device(config)# no ip pimsm-snooping
```

If you also want to disable IP multicast traffic reduction, enter the following command.

```
device(config)# no ip multicast
```

Multicast traffic reduction per VLAN or VPLS instance

You can configure the following methods for reducing multicast traffic globally on a Extreme device:

- IGMP snooping - This is described in [Changing the IGMP mode](#) on page 14.
- PIM snooping - This is described in [PIM SM traffic snooping](#) on page 16.

When these are set globally on a router, they apply to all VLANs and all VPLS instances that are configured on the router. You can configure specified VLANs or VPLS instances for multicast traffic reduction by these methods as described in the following sections. Additionally, you are able to configure IGMP and PIM proxy which are only configurable per VLAN or VPLS instance.

Multicast traffic reduction per VPLS instance is supported for dual-mode, untagged, single-tagged, and dual-tagged VPLS endpoints.


Configuration Notes

- IP multicast traffic reduction per VPLS instance supports a maximum of up to 2000 IP multicast groups.
- IGMP snooping cannot be concurrently enabled on MLX Series and XMR Series devices with VPLS CPU Protection on a VPLS instance.
- IGMP Snooping cannot be configured on MLX Series and XMR Series devices when a VPLS instance has ISID endpoints.
- Traffic will continue to be forwarded only to those VPLS endpoints or peers from which a join for the (S, G) has been received irrespective of the status of the local switching option.

Application example

Figure 3 shows an example of multicast traffic reduction in a VPLS network.

FIGURE 3 IP multicast traffic reduction in a VPLS network



In the example shown in [Figure 3](#), when IP multicast traffic reduction (IGMP snooping) is enabled on the VPLS network, PE 1 will be selected as the active port (querier) because it has the lowest router ID amongst the PEs in the VPLS instance. PE 1 will actively send out IGMP queries to solicit information from IP multicast groups within the VPLS instance. PE 2 and PE 3 will not send out queries as they are in passive mode, but will respond to the query from PE 1, with the respective report information received from their hosts.

In a snooping configuration within a VPLS instance, multicast traffic is always flooded to the device that is in active IGMP mode (the router port). If the IP multicast domain includes an IP multicast router, that router will be the querier. In this case, the querier is also known as the router port. If there is no IP multicast router in the domain, one of the devices in the network can be manually configured for active IGMP mode. Otherwise, the device with the lowest router ID will be elected as the querier. In the example in [Figure 3](#), PE 1 is the elected querier.

In a VPLS scenario, reports are always forwarded to all VPLS peers whether or not the peer is the querier. In the above example, PE 1 is elected the querier, but if PE 2 is connected to a receiver, it forwards the reports received from the receiver to both VPLS peers PE 3 and PE 1. Therefore, any traffic for the host connected to PE 2 received from PE 3 are sent by PE 3 to both the router port PE 1 as well as the receiver PE 2.

PE 1 drops the traffic received from PE 3 because it is aware of the presence of a receiver attached to the router PE 2. Traffic sent from PE 3 will only be received by the host connected to PE 2.

If there is no receiver present in the setup, then traffic from PE 3 will only be flooded to its local endpoints and the router port PE 1. Router PE 1 sends the traffic out of its endpoints.

In case of PIM SM snooping, traffic received for unknown groups is always dropped. There is no router port concept in case of PIM SM snooping.

Multicast Traffic Forwarding

MLX Series and XMR Series devices use the IP multicast group address for data forwarding when supporting multicast traffic reduction over VPLS.

On CES 2000 Series and CER 2000 Series devices, data forwarding uses the destination multicast MAC address. The IP multicast group address is mapped to its destination MAC address. The MAC address is programmed as the VPLS MAC entry for the IP multicast group.

Implementing data forwarding, using the destination multicast MAC address, has the following limitations on CES 2000 Series and CER 2000 Series devices:

- Data forwarding based on the source IP address is not supported.
- All packets with destination multicast MAC address are forwarded, including IP unicast packets, IP multicast packets that did not follow the standard mapping, and non-IP packets.
- When a number of IP group addresses map to the same destination multicast MAC address, the output ports are the superset of all the output ports of the IP group addresses mapped to the destination multicast MAC address. As a result, some hosts receive unwanted traffic. For example, IP groups G1 and G2 map to the same destination multicast MAC address. If G1 contains port 1 and port 2 and G2 contains port 2 and port 3, traffic sent to G1 and G2 will be forwarded to ports 1, 2, and 3. Hosts that are connected to port 1 and port 3 will receive unwanted traffic.

Configuring the IGMP mode per VLAN or VPLS instance

In the following example, multicast traffic reduction is applied using IGMP snooping to VLAN 2.

```
device(config)# vlan 2
device(config-vlan-2)# multicast passive
```

To remove multicast traffic reduction configurations in VLAN 2, and take the global multicast traffic reduction configuration, enter the following command.

```
device(config)# vlan 2
device(config-vlan-2)# no multicast
```

In the following example, multicast traffic reduction is applied using IGMP snooping to VPLS instance V1.

```
device(config)# router mpls
device(config-mpls)# vpls v1 10
device(config--mpls-vpls-v1)# multicast passive
```

Syntax: [no] multicast active | passive

When you enable IP multicast for a specific VLAN or VPLS instance, IGMP snooping is enabled. The device uses IGMP to maintain a table of the Group Membership reports received by the device for the specified VLAN or VPLS instance. You can use active or passive IGMP mode. There is no default mode.

The description for the IGMP modes is as follows:

- **Active** - When active IGMP mode is enabled, the router actively sends out IGMP queries to identify IP multicast groups within the VLAN or VPLS instance and makes entries in the IGMP table based on the Group Membership reports received from the network.
- **Passive** - When passive IGMP mode is enabled, the router listens for IGMP Group Membership reports on the VLAN or VPLS instance specified but does not send IGMP queries. The passive mode is called "IGMP snooping". Use this mode when another device in the VLAN or VPLS instance is actively sending queries.

Configuring the PIM SM traffic snooping per VLAN or VPLS instance

In the following example, multicast traffic reduction is applied using PIM SM Traffic snooping to VLAN 2.

```
device(config)# vlan 2
device(config-vlan-2)# multicast pimsm-snooping
```

In the following example, multicast traffic reduction is applied using PIM SM traffic snooping to VPLS instance V1.

```
device(config)# router mpls
device(config-mpls)# vpls v1 10
device(config--mpls-vpls-v1)# multicast pimsm-snooping
```

Syntax: [no] multicast pimsm-snooping

Configuring PIM proxy per VLAN or VPLS instance

Using the PIM proxy function, multicast traffic can be reduced by configuring an Extreme device to issue PIM join and prune messages on behalf of hosts that the configured router discovers through standard PIM interfaces. The router is then able to act as a proxy for the discovered hosts and perform PIM tasks upstream of the discovered hosts. Where there are multiple PIM downstream routers, this removes the need to send multiple messages.

To configure a Extreme device to function as a PIM proxy on VLAN 2, use the following commands.

```
device(config)# vlan 2
device(config-vlan-2)# multicast pim-proxy-enable
```

To configure a Extreme device to function as an PIM proxy on VPLS instance V1, use the following commands.

```
device(config)# router mpls
device(config-mpls)# vpls v1 10
device(config--mpls-vpls-v1)# multicast pim-proxy-enable
```

Syntax: [no] multicast pim-proxy-enable

Configuring IGMP snooping tracking per VLAN or VPLS instance

When IGMP Snooping Tracking is enabled, the Extreme device immediately removes any IGMP host port from the IP multicast group entry when it detects an IGMP-leave message on the specified host port without first sending out group-specific queries to the interface. By default, IGMP Snooping Tracking is disabled.

The **ip multicast tracking** command may be enabled globally as well as per VLAN basis. To enable IGMP Snooping Tracking globally, enter a command such as the following.

```
device(config)# multicast tracking
```

Syntax: [no] ip multicast tracking

The **no** form of this command disables the tracking process globally.

To enable IGMP Snooping Tracking per VLAN, enter commands such as the following.

```
device(config)# vlan 100
device(config-vlan-100)# multicast tracking
```

Syntax: [no] multicast tracking

The **no** form of this command disables the tracking process per VLAN.

To enable IGMP Snooping Tracking per VPLS instance, enter commands such as the following.

```
device(config)# router mpls
device(config-mpls)# vpls vl 10
device(config--mpls-vpls-vl)# multicast tracking
```

For IGMPv3, the above command also internally tracks all the IGMPv3 hosts behind a given port. The port is not removed from the IP multicast group entry in the forwarding table until all the hosts behind that port have left that multicast group. When the last IGMPv3 host sends a IGMPv3 leave message, the port is removed from the IP multicast group entry in the forwarding table immediately without first sending out group_source_specific query to the interface

Syntax: [no] multicast tracking

The **no** form of this command disables the tracking process per VPLS instance.

Multicast snooping over VPLS will not load-balance the multicast traffic among multiple tunnels,

Multicast snooping over VPLS will not load-balance the multicast traffic among multiple tunnels,

NOTE

Multicast snooping over VPLS will not load-balance the multicast traffic among multiple tunnels when IGMP Snooping is enabled.

IP Multicast CPU Protection

When IGMP snooping is enabled, IP Multicast CPU Protection protects the MP CPU from high CPU usage by the multicast tasks in scenarios where there is a large number of streams for specific groups with large number of OIFs thereby resulting in numerous pending FID updates.

IP Multicast CPU Protection introduces (*,G) based forwarding for snooping to relieve the MP CPU of numerous pending FID updates, and aging out OIFs, especially in cases when there are multiple sources and receivers for a specific group G. The ACL modification and deletion events are reflected in the CPU protection entries.

IP Multicast CPU protection also handles the addition of new VLANs. They inherit the global CPU protection configuration if any.

NOTE

Before configuring or adding any other clause with new group addresses, you should check for matching S, G entries in the existing access-list. Configure or add the new clause(s) to the existing S, G entries as appropriate. Create new S, G entries only if there is no match. You will optimize the use of multicast resources by saving room for adding new groups since the maximum number of groups you can configure or add is limited.

ACL traffic filtering

An extended ACL needs to be provided as the parameter. It should have the group(s) for which cpu-protection is intended. As such, the ACL must contain statement(s) equivalent to the form **permit ip|ipv6 any any G/32** where G is the multicast group for which cpu-protection is intended.

The ACL provided as parameter may be already defined or defined later after the initial configuration. Whenever there is a change in that ACL, previous cpu-protection groups will be deleted and reprogrammed with the cpu-protection group(s) for relevant VLANs. An undefined ACL or ACLs without having clauses equivalent to the form **permit ip|ipv6 any any G/32** will not be programming any cpu-protection entries.

Configuration Considerations

- The **ip multicast cpu-protection** command also supports IPv6 traffic. The **ipv6 multicast cpu-protectionacl-name** command must be configured under the IPv6 router configuration to support IPv6. Only ACL names are supported when using this command with IPv6.
- CPU protection can be configured both at the global level as well as vlan-specific level.
- Global CPU protection configuration is inherited by all current and new VLAN.
- An extended ACL needs to be provided as the parameter. It should have the group(s) for which cpu-protection is intended. As such, the ACL must contain statement(s) equivalent to the form **permit ip|ipv6 any any G/32** where G is the multicast group for which cpu-protection is intended.
- The ACL provided as parameter may be already defined or defined later after configuration. Whenever there is a change in that ACL, previous cpu-protection groups will be deleted and reprogrammed with the cpu-protection group(s) for relevant VLANs. An undefined ACL or ACLs without having clauses equivalent to the form **permit ip|ipv6 any any G/32** will not be programming any cpu-protection entries.

Enabling IP multicast CPU Protection

To globally enable IP Multicast CPU Protection, enter the following command.

```
device(config)# ip multicast cpu-protection 101
```

Syntax: **[no] ip multicast cpu-protection acl-id or acl-name**

The *acl-id* or *acl-name* parameter specifies the ACL ID or name used to configure CPU protection.

The **no** option removes the CPU protection configuration. Only one ACL name or ACL ID can be applied. The cpu-protection ACL name or ACL ID should match with the one applied to the configuration.

To globally enable IPv6 Multicast CPU Protection, enter the following command.

```
device(config)# ipv6 multicast cpu-protection 101
```

The *acl-name* parameter specifies the ACL name used to configure CPU protection.

The **no** option removes the CPU protection configuration. Only one ACL name or ACL ID can be applied. The cpu-protection ACL name or ACL ID should match with the one applied to the configuration.

Syntax: [no] ipv6 multicast cpu-protection *acl-name*

Enabling multicast CPU Protection on a VLAN

To enable multicast CPU protection on a VLAN, enter the following command.

```
device(config-vlan-10)# multicast cpu-protection 101
```

Syntax: [no] ip multicast cpu-protection

To globally enable IPv6 Multicast CPU Protection, enter the following command.

```
device(config-vlan-10)# ipv6 multicast cpu-protection 101
```

Syntax: [no] ipv6 multicast cpu-protection *acl-name*

The *acl-id* or *acl-name* parameter specifies the ACL ID or name used to configure CPU protection.

The **no** option removes the CPU protection configuration. Only one ACL name or ACL ID can be applied. The cpu-protection ACL name or ACL ID should match with the one applied to the configuration.

Verifying CPU protection

The **multicast cpu-protection** global level command configures itself implicitly on all the VLANs.

```
device show ip multicast vlan 18
-----+-----+-----+-----+-----+-----+-----+-----
VLAN State Mode      Active      Time (*, G) (S, G)
              Querier      Query Count Count
-----+-----+-----+-----+-----+-----+-----+-----
18  Ena  Active  Self      96      3      2
-----+-----+-----+-----+-----+-----+-----+-----
Router ports:
Flags: R-Router Port, V2|V3: IGMP Receiver, P_G|P_SG: PIM Join
  1    (*, 229.15.15.1 ) 00:19:15 NumOIF: 3 profile: 31
      Outgoing Interfaces:
          e3/3 vlan 18 ( V2) 00:17:28/22s
          e1/9 vlan 18 ( V2) 00:19:09/21s
          e1/24 vlan 18 ( V2) 00:19:09/23s
      FID: 0xa28e      MVID: None
      Cam_idx: 0x0005cf80 (PPCR-2)    04:09:25
          Cam_idx: 0x0005d268 (PPCR-1)    04:09:25
```

Syntax: show ip multicast vlan *vlan-id*

The number of cam-indices allocated for the entry on an LP depends upon the number of PPCRs on which the VLAN has member ports.

The **vlan** *vlan-id* parameter displays IP multicast VLAN information for a specified VLAN.

[Table 2](#) describes the output parameters of the **show ip multicast vlan** command.

TABLE 2 Output parameters of the show ip multicast vlan command

Field	Description
VLAN	Shows the ID of the configured VLAN.
State	Shows whether the VLAN interface is enabled or disabled.

TABLE 2 Output parameters of the show ip multicast vlancommand (continued)

Field	Description
Mode	Shows whether the VLAN interface is in active mode or passive mode.
Active Querier	Shows the active IGMP querier for the VLAN.
Time Query	Shows the time countdown to generate the next query message.
(* ,G)Count	Shows the count of (*,G) entries.
(S, G)Count	Shows the count of (S,G) entries.
Flags	Shows the flags of the outgoing interface.
V2 V3	Shows the version of the IGMP message received.
P_G	Indicates that a PIM (*,G) join was received on that interface.
P_SG	Indicates that a PIM (S,G) join was received on that interface.
NumOIF	Show the count of the outgoing interface.
profile	Shows the profile ID associated with the stream.
Outgoing Interfaces	Shows the list of outgoing interfaces.
FID	Shows the FID resource allocated for a particular entry.
MVID	Shows the MVID resource allocated for a particular entry.

Static IGMP membership

When configuring a static IGMP membership, you have two options:

The **multicast static-group uplink** command which sends the traffic to the router, and saves a port.

The **multicast static-group group-address port-list** command is for downstream traffic and uses a port.

Configuring a multicast static group uplink per VLAN

When the **multicast static-group uplink** command is enabled on a snooping VLAN, the snooping device behaves like an IGMP host on ports connected to the multicast router. The snooping device will respond to IGMP queries from the uplink multicast PIM router for the groups and sources configured. Upon the multicast router receiving the IGMP join message, it will initiate the PIM join on its upstream path towards the source to pull the source traffic down. The source traffic will stop at the IGMP snooping device. The traffic will then be forwarded to the multicast receiver and router ports or dropped in hardware if no other multicast receiver and routers are present in the VLAN.

The **multicast static-group uplink** command cannot be configured globally per VPLS basis. It can be configured under the VLAN configuration only.

The **multicast static-group uplink** command must be used with the **multicast static-group** command in order to connect a remote multicast source with the snooping vlan where the static-group is configured.

When using IGMP v3, you can use the **multicast static-group include** or **multicast static-group exclude** command to statically **include** or **exclude** multicast traffic, respectively for hosts that cannot signal group membership dynamically.

To configure the snooping device to statically join a multicast group on the uplink interface, enter commands such as the following.

```
device(config)# vlan 100
device(config-vlan-100)# multicast static-group 224.10.1.1 uplink
```

To configure the physical interface 10.43.3.12 to statically join a multicast group on port 2/4, enter commands such as the following.

```
device(config)# vlan 100
device((config-vlan-100)# multicast static-group 224.10.1.1 2/4
```

To configure the snooping device to statically join a multicast stream with the source address of 10.43.1.12 in the include mode, enter commands such as the following.

```
device(config)# vlan 100
device(config-vlan-100)# multicast static-group 224.10.1.1 include 10.43.1.12 uplink
```

To configure the snooping device to statically join all multicast streams on the uplink interface excluding the stream with source address 10.43.1.12, enter commands such as the following.

```
device(config)# vlan 100
device(config-vlan-100)# multicast static-group 224.10.1.1 exclude 10.43.1.12 uplink
```

Configuring multicast static group port-list per VLAN

When the **multicast static-group group-address port-list** command is enabled on a snooping VLAN, the snooping device will add the ports to the outgoing interface list of the multicast group entry in the forwarding table as if IGMP joins were received from these ports. These ports will not be aged out from the multicast group for not responding to the IGMP queries.

The **multicast static-group group-address port-list** command cannot be configured globally per VPLS basis.

It can be configured under the VLAN configuration level only.

To configure the physical interface ethernet 2/4 to statically join a multicast group, enter commands such as the following.

```
device(config)# vlan 100
device(config-vlan-100)# multicast static-group 224.10.1.1 ethernet 2/4
```

To configure the physical interface ethernet 3/4 to statically join a multicast stream with source address of 10.43.1.12 in the include mode, enter commands such as the following.

```
device(config)# vlan 100
device(config-vlan-100)# multicast static-group 224.10.1.1 include 10.43.1.12 ethernet 3/4
```

To configure the physical interface ethernet 3/4 to statically join all multicast streams on the uplink interface excluding the stream with source address of 10.43.1.12, enter commands such as the following.

```
device(config)# vlan 100
device(config-vlan-100)# multicast static-group 224.10.1.1 exclude 10.43.1.12 ethernet 3/4
```

Syntax: [no] **multicast static-group** *group-address* **uplink**

Syntax: [no] **multicast static-group** *group-address* **port-list**

IGMP v3 Commands

Syntax: [no] **multicast static-group** *group-address* [**include** | **exclude** *source-address*] **uplink**

Syntax: [no] **multicast static-group** *group-address* [**include** | **exclude** *source-address*] **port-list**

The **group-address** parameter specifies the group multicast address.

The **include** or **exclude** keyword indicates a filtering action. You can specify which source (for a group) to include or exclude. The include or exclude keyword is only supported on IGMPv3.

The **source-address** parameter specifies the IP address of the multicast source. Each address must be added or deleted one line per source.

The **uplink** parameter specifies the port as an uplink port that can receive multicast data for the configured multicast groups. Upstream traffic will be sent to the router and will not use a port.

The **port-list** parameter specifies the range of ports to include in the configuration.

The **no** form of this command removes the static multicast definition. Each configuration must be deleted separately.

Displaying IP multicast information

The following sections show how to display and clear IP multicast reduction information.

Displaying multicast information

You can display IP multicast traffic information in a brief form for all instances or in a detail form for a specified VLAN or VPLS instance.

The following example shows statistics for a specific VLAN.

The output of this command now displays a separate timer for each flag type (*, G) entry. The multicast snooping per flag aging enhancement provides support to maintain aging of flags per OIF for a snooping entry.

```
device# show ip multicast vlan 1500
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
VLAN      State Mode      Active      Time (*, G) (S, G)
Querier   Query Count Count
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1500      I-Ena Passive  10.25.10.10  103   1   3
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

Router ports: 7/16 (60s)

Flags- R: Router Port,  V2|V3: IGMP Receiver,  P_G|P_SG: PIM Join

  1   (*, 239.10.10.10) Uptime: 00:01:38      NumOIF: 1      profile: none
      Outgoing Interfaces:
          e7/16 vlan 1500  00:01:38 Flags: ( R [60s] V2 [72s])

  1   (10.25.120.131, 239.10.10.10) in e4/1 vlan 1500 Uptime: 00:00:06      NumOIF: 1  profile: none
      Outgoing Interfaces:
          e7/16 vlan 1500  00:00:06 Flags: ( R V2)

      FID: 0xa013      MVID: None

  2   (10.25.120.130, 239.10.10.10) in e4/1 vlan 1500 Uptime: 00:00:06      NumOIF: 1  profile: none
      Outgoing Interfaces:
          e7/16 vlan 1500  00:00:06 Flags: ( R V2)

      FID: 0xa012      MVID: None

  3   (10.25.120.129, 239.10.10.10) in e4/1 vlan 1500 Uptime: 00:01:39      NumOIF: 1  profile: none
      Outgoing Interfaces:
          e7/16 vlan 1500  00:01:39 Flags: ( R V2)

      FID: 0xa011      MVID: None
```

To display detailed IP multicast traffic reduction information for a specified VPLS instance on the Brocade devices, enter the **show ip multicast vpls** command at any level of the CLI:

```
device# show ip multicast vpls 33
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
VPLS      State Mode      Active      Time (*, G) (S, G)
Querier   Query Count Count
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
33        I-Ena Passive  10.174.63.254  15   400  400
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

Router ports: TNNL peer 10.10.1.1 (22s) VC Label 983114

Flags- R: Router Port,  V2|V3: IGMP Receiver,  P_G|P_SG: PIM Join
```

```

1    (*, 233.17.169.97) 09:05:18      NumOIF: 2      profile: none
    Mapped MAC Address: 0100.5e11.a961 FID: 0x05b0
    Outgoing Interfaces:
      e1/3 vlan 33 09:05:18 Flags: ( V2 [60s] V3 [72s])
      TNNL peer 10.10.1.1 09:05:18 Flags: ( V2 [22s])

1    (10.150.104.1, 233.17.169.97) in TNNL peer 10.10.1.1 09:05:18      NumOIF: 2      profile: none
    Outgoing Interfaces:
      e1/3 vlan 33 09:05:18 Flags: ( V2 V3)
      TNNL peer 10.10.1.1 09:05:18 Flags: ( V2)
FID: 0x02df      MVID:      0

```

NOTE

The italicized text in the output of the command **show ip multicast vpls vpls-id** command indicates multicast MAC address information and is displayed only in the output for Brocade NetIron CES and Brocade NetIron CER devices. Multicast MAC address information is not displayed in the output of Brocade MLX series and Brocade NetIron XMR devices.

To display only multicast MAC address information for Brocade NetIron CES and Brocade NetIron CER devices, enter the **show ip multicast vpls mac** command at any level of the CLI.

NOTE

The **show ip multicast vpls mac** and **show ip multicast vpls vpls-id mac** commands are not supported on Brocade MLX series and Brocade NetIron XMR devices.

You can display PIM SM information by entering the **show ip multicast vlan vlan-id pim** command at any level of the CLI.

```

device(config)# show ip multicast vlan 2 pim
Number of PIM Groups:5
Total Number of PIM Entries: 5
vlan group      port join-source age prune-source age
-----
2    229.0.0.5    1/1 2.1.1.100    180
2    229.0.0.4    1/1 2.1.1.100    180
2    229.0.0.3    1/1 2.1.1.100    180
2    229.0.0.2    1/1 2.1.1.100    180
2    229.0.0.1    1/1 2.1.1.100    180
device(config)#

```

Displaying IP multicast statistics

To display IP multicast statistics on a device, enter the **show ip multicast vlan vlan-id statistics** command at any level of the CLI.

```

device# show ip multicast vlan 1 statistics
IP multicast is enabled - Passive
VLAN ID 1
Reports Received:          34
Leaves Received:          21
General Queries Received: 60
Group Specific Queries Received: 2
Others Received:           0
General Queries Sent:      0
Group Specific Queries Sent: 0

```

The command in the example shows the statistics for VLAN 1.

Syntax: **show ip multicast vlan vlan-id statistics**

Clearing IP multicast statistics

To clear IP multicast statistics on a device, enter the following command at the Privileged EXEC level of the CLI.

```
device# clear ip multicast statistics
```

This command resets statistics counters for all the VLANs to zero.

Syntax: `clear ip multicast statistics`

Clearing IGMP group flows

To clear all the IGMP flows learned by the device, enter the following command at the Privileged EXEC level of the CLI.

```
device# clear ip multicast all
```

The following example shows IGMP flows information listed by the **show ip multicast** command, followed by removal of the information by the **clear ip multicast all** command.

```
device# show ip multicast
IP multicast is enabled - Active
VLAN ID 1
Active 192.168.2.30 Router Ports 4/13
Multicast Group: 239.255.162.5, Port: 4/4 4/13
Multicast Group: 239.255.162.4, Port: 4/10 4/13
device# clear ip multicast all
device# show ip multicast
IP multicast is enabled - Active
VLAN ID 1
Active 192.168.2.30 Router Ports 4/13
```

To clear the learned IGMP flows for a specific IP multicast group, enter a command such as the following.

```
device# clear ip multicast vlan 1 group 239.255.162.5
```

NOTE

Layer 2 IGMP snooping S,G entries do not time out on their own on Netron CER devices. You must use the **clear ip multicast vlan-id group group-id** command to clear specific entries.

The following example shows how to clear the IGMP flows for a specific group and retain reports for other groups.

```
device# show ip multicast
IP multicast is enabled - Active
VLAN ID 1
Active 192.168.2.30 Router Ports 4/13
Multicast Group: 239.255.162.5, Port: 4/4 4/13
Multicast Group: 239.255.162.4, Port: 4/10 4/13
device# clear ip multicast vlan 1 group 239.255.162.5
device# show ip multicast
IP multicast is enabled - Active
VLAN ID 1
Active 192.168.2.30 Router Ports 4/13
Multicast Group: 239.255.162.4, Port: 4/10 4/13
```

Syntax: `clear ip multicast vlan vlan-id all | group group-id`

The **all** parameter clears the learned flows for all groups.

The **group** group-id parameter clears the flows for the specified group but does not clear the flows for other groups.

IPv6 Multicast VLAN Traffic Reduction

- IPv6 Multicast Listener Discovery snooping..... 31

IPv6 Multicast Listener Discovery snooping

IPv6 Multicast Listener Discovery (MLD) snooping controls the amount of multicast traffic in a switched network. By default, a LAN switch floods the broadcast domain with multicast IPv6 packets. If many multicast servers are sending streams to the segment, this will consume a lot of bandwidth. MLD snooping identifies multicast-enabled router ports and multicast receiver ports in a given VLAN or a switched network and forwards multicast traffic only to those ports.

Configuring IPv6 multicast routing or snooping

IPv6 multicast snooping or routing can be enabled on a VE interface or VLAN, but not on both. This is because all of the multicast data and control packets received on the snooping VLAN are handled by multicast snooping and do not reach the multicast routing component. Similarly, any multicast data or control packets received on a VE interface enabled with PIM routing are handled by the PIM routing component and are not seen by the MLD snooping component.

The following considerations apply when configuring concurrent operation of multicast routing and snooping.

- Either multicast snooping or routing can be enabled on a VE or VLAN but not both.
- Snooping can be enabled globally **ipv6 multicast active | passive**.
- The global snooping configuration is inherited by all current VLANs that are not enabled for multicast routing.
- The global snooping configuration is also inherited by all new VLANs. To enable multicast routing on a newly configured VE or VLAN (when snooping is globally enabled), you must first disable snooping on the newly created VE or VLAN.
- Global snooping configuration must be configured first before VLAN configuration.
- A VLAN-level snooping configuration is displayed only if it is different from the global configuration.

NOTE

On a snooping switch, IPv6 enable must be configured on the VE or VLAN that has MLD snooping enabled.

Enabling IPv6 multicast traffic reduction

By default, the device forwards all IPv6 multicast traffic out to all ports except the port on which the traffic was received. To reduce multicast traffic through the device, you can enable IPv6 Multicast Traffic Reduction. This feature configures the device to forward multicast traffic only on the ports attached to multicast group members, instead of forwarding all multicast traffic to all ports. The device determines the ports that are attached to multicast group members based on entries in the MLD Snooping table. Each entry in the table consists of MAC addresses and the ports from which the device has received Group Membership reports for that group.

By default, the device broadcasts traffic addressed to an IPv6 multicast group that does not have any entries in the MLD Snooping table. When you enable IPv6 Multicast Traffic Reduction, the device determines the ports that are attached to multicast group members based on entries in the MLD Snooping table. The MLD Snooping table entries are created when the VLAN receives a Group Membership report for a group. Each entry in the table consists of an IPv6 multicast group address and the ports from which the device has received Group Membership reports.

When the device receives traffic for an IPv6 multicast group, the device looks in the MLD Snooping table for an entry corresponding to that group. If the device finds an entry, the device forwards the group traffic out to the ports listed in the corresponding entries, as long as the ports are members of the same VLAN. If the table does not contain an entry corresponding to the group or if the port is a member of the default VLAN, the device broadcasts the traffic.

When one or more devices are running Layer 2 IPv6 Multicast Traffic Reduction, configure one of the devices for active MLD and leave the other devices configured for passive MLD. However, if the IPv6 multicast domain contains a multicast-capable router, configure all the devices for MLD and allow the router to actively send the MLD queries.

Configuring IPv6 MLD snooping

To enable IPv6 Multicast Traffic Reduction, enter the following command.

```
device(config)# ipv6 multicast active
```

Syntax: [no] ipv6 multicast active | passive

When you enable IPv6 multicast on a device, all ports on the device are configured for MLD.

If you are using passive MLD, all ports can send MLD queries and receive MLD reports. If you are using active MLD, all ports can receive MLD queries.

IPv6 Multicast Traffic Reduction cannot be disabled on individual ports of a device. IPv6 Multicast Traffic Reduction can be disabled globally by entering the **no ipv6 multicast** command.

To verify that IPv6 Multicast Traffic Reduction is enabled, enter the following command at any level of the CLI.

```
device(config)# show ipv6 multicast
IPv6 multicast is enabled - Active
```

Syntax: show ipv6 multicast

Changing the MLD mode

When you enable IPv6 Multicast Traffic Reduction on the device, MLD also is enabled. The device uses MLD to maintain a table of the Group Membership reports received by the device. You can use active or passive MLD mode. There is no default mode.

The active and passive MLD modes are described as follows:

- **Active** - When active MLD mode is enabled, the device actively sends out MLD queries to identify IPv6 multicast groups on the network and makes entries in the MLD table based on the Group Membership reports received from the network.

NOTE

Routers in the network generally handle this operation. Use the active MLD mode only when the device is in a standalone Layer 2 switched network with no external IPv6 multicast router attachments. In this case, enable the active MLD mode on only one of the devices and leave the other devices configured for passive MLD mode.

- **Passive** - When passive MLD mode is enabled, the device listens for MLD Group Membership reports but does not send MLD queries. The passive mode is sometimes called "MLD snooping". Use this mode when another device in the network is actively sending queries.

Globally configuring the age interval

When the device receives a Group Membership report, the device makes an entry in the MLD group table for the group in the report. The age interval specifies how long the entry can remain in the table without the device receiving another Group Membership report.

To configure the age interval, enter a command such as the following.

```
device(config)# ipv6 multicast age-interval 280
```

Syntax: `[no] ipv6 multicast age-interval interval`

The *interval* parameter specifies the interval between queries. You can specify a value from 10 through 1220 seconds. The default is 260 seconds.

Filtering multicast groups

By default, the device forwards multicast traffic for all valid multicast groups. You can configure a device to filter out all multicast traffic for groups other than the ones for which the device has received Group Membership reports.

When the device starts up, it forwards all multicast groups even though multicast traffic filters are configured. This process continues until the device receives a Group Membership report. Once the Group Membership report is received, the device drops all multicast packets for groups other than the ones for which the device has received the Group Membership report.

To enable IPv6 multicast filtering, enter the following command.

```
device(config)# ipv6 multicast filter
```

Syntax: `[no] ipv6 multicast filter`

Setting PIM proxy interval

NOTE

The PIM proxy interval value should not be changed. Changing the PIM proxy interval value is not supported.

The PIM proxy interval specifies the time interval in seconds between the PIM proxy and join messages.

To set the time interval between PIM proxy messages, enter the following command.

```
device(config)# ipv6 multicast pim-proxy-interval 10
```

Syntax: `ipv6 multicast pim-proxy-interval interval`

The *interval* parameter specifies the interval between PIM proxy messages. You can specify a value from 10 through 600 seconds.

PIM-SM traffic snooping

By default, when a device receives an IPv6 multicast packet, the device does not examine the multicast information in the packet. Instead, the device simply forwards the packet out to all ports except the port that received the packet. In some networks, this method can cause unnecessary traffic overhead in the network. For example, if the device is attached to only one group source and two group receivers, but has devices attached to every port, the device forwards group traffic out all ports in the same broadcast domain except the port attached to the source, even though there are only two receivers for the group.

PIM-SM traffic snooping eliminates the superfluous traffic by configuring the device to forward IPv6 multicast group traffic only on the ports that are attached to receivers for the group.

PIM-SM traffic snooping requires IPv6 Multicast Traffic Reduction to be enabled on the device. IPv6 Multicast Traffic Reduction configures the device to listen for MLD messages. PIM-SM traffic snooping provides a finer level of multicast traffic control by configuring the device to listen specifically for PIM-SM join and prune messages sent from one PIM-SM router to another through the device.

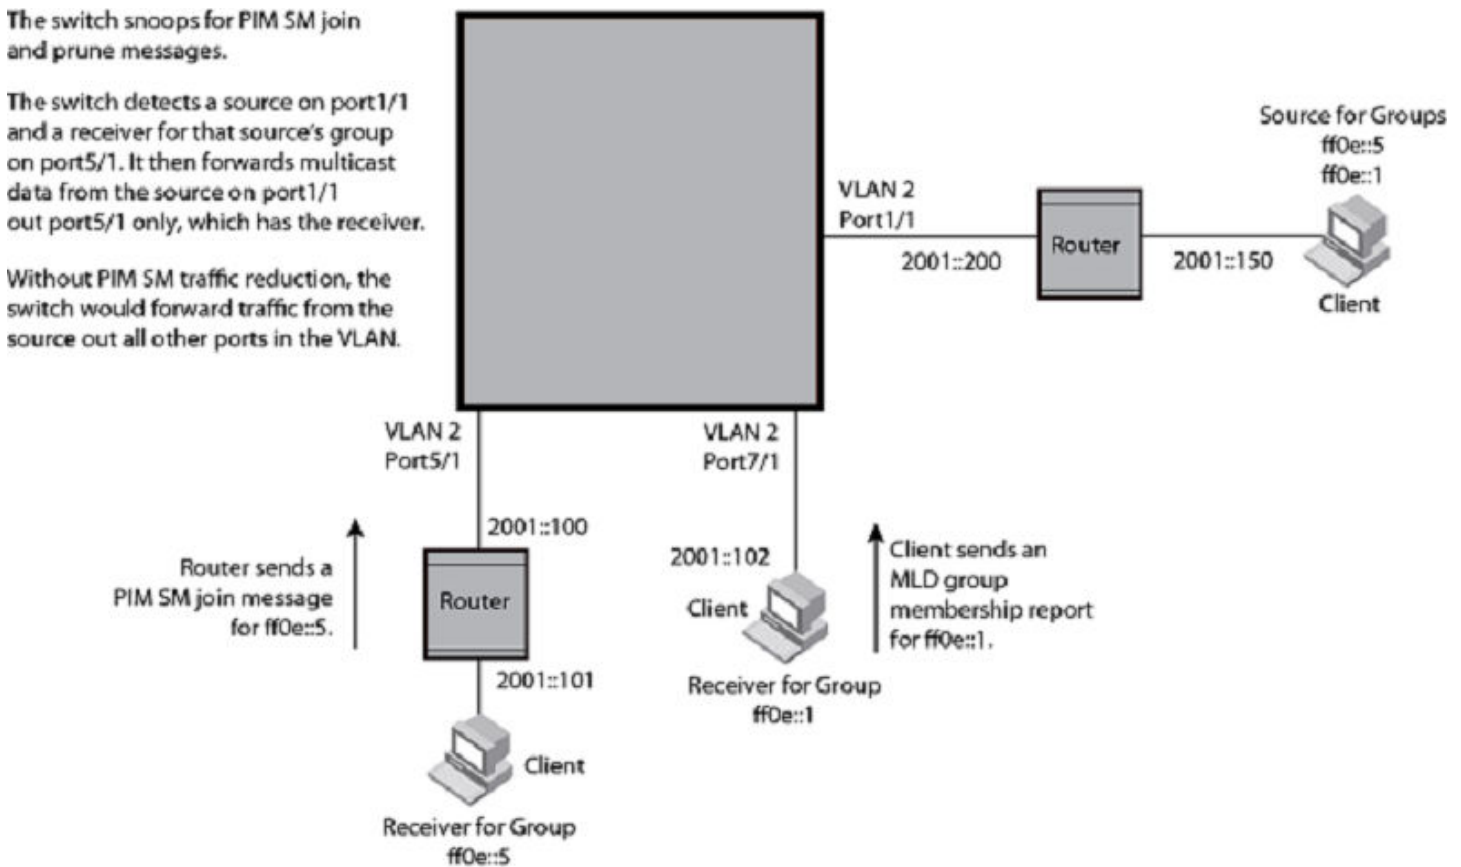
Application examples

Figure 4 shows an example application of the PIM-SM traffic snooping feature. In this example, a device is connected through an IP router to a PIM-SM group source that is sending traffic for two PIM-SM groups. The device also is connected to a receiver for each of the groups.

NOTE

PIM-SM traffic snooping applies only to PIM-SM version 2 (PIM V2).

FIGURE 4 PIM-SM IPv6 traffic reduction in enterprise network



When PIM-SM traffic snooping is enabled, the device starts listening for PIM-SM join and prune messages and MLD Snooping reports. Until the device receives a PIM-SM join message or an MLD report, the device forwards IPv6 multicast traffic out to all ports. Once the device receives a join message or Group Membership report for a group, the device forwards subsequent traffic for that group only on the ports from which the join messages or MLD reports were received.

In this example, the router connected to the receiver for group ff0e::1 sends a join message toward the source of the group. Since PIM-SM traffic snooping is enabled on the device, the device examines the join message to learn the group ID, then makes a forwarding entry for the group ID and the port connected to the router of the receiver. The next time the device receives traffic for ff0e::1 from the source of the group, the device forwards the traffic only on port 5/1, because that is the only port connected to a receiver for the group.

Notice that the receiver for group ff0e::5 is directly connected to the device. As a result, the device does not see a join message on behalf of the client. However, because IP Multicast Traffic Reduction also is enabled, the device uses the MLD Group Membership report from the client to select the port for forwarding traffic to group ff0e::5 receivers.

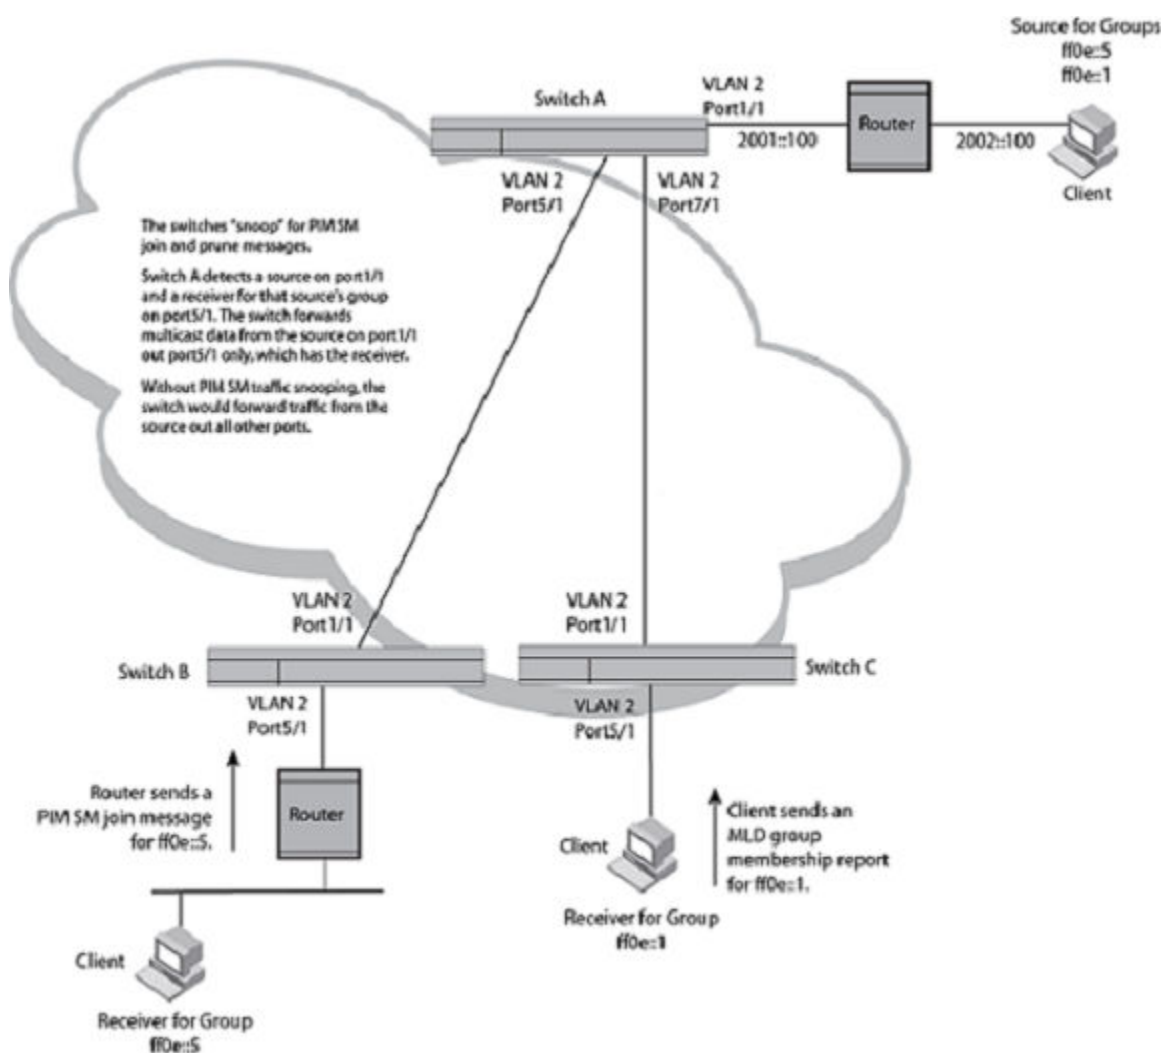
The IPv6 Multicast Traffic Reduction feature and the PIM-SM traffic snooping feature together build a list of groups and forwarding ports for the VLAN. The list includes PIM-SM groups learned through join messages as well as MAC addresses learned through MLD reports. In this case, even though the device never sees a join message for the receiver for group ff0e::5, the device nonetheless learns about the receiver and forwards group traffic to the receiver.

The device stops forwarding IPv6 multicast traffic on a port for a group if the port receives a prune message for the group.

Notice that the ports connected to the source and the receivers are all in the same port-based VLAN on the device. This is required for the PIM-SM traffic snooping feature. The feature also requires the source and the downstream router to be on different IP subnets, as shown in Figure 4.

Figure 5 shows another example application for PIM-SM traffic snooping. This example shows devices on the edge of a global Ethernet cloud. Assume that each device is attached to numerous other devices.

FIGURE 5 PIM-SM IPv6 traffic reduction in global Ethernet environment



The devices on the edge of the global Ethernet cloud are configured for IP Multicast Traffic Reduction and PIM-SM traffic snooping. Although this application uses multiple devices, the feature has the same requirements and works the same way as it does on a single device.

Configuration requirements

Consider the following configuration requirements:

- IPv6 Multicast Traffic Reduction must be enabled on the device that will be running PIM-SM snooping. The PIM-SM traffic snooping feature requires IPv6 Multicast Traffic Reduction.

NOTE

Use the passive mode of IPv6 Multicast Traffic Reduction instead of the active mode. The passive mode assumes that a router is sending group membership queries as well as join and prune messages on behalf of receivers. The active mode configures the device to send group membership queries.

- All the device ports connected to the source and receivers or routers must be in the same port-based VLAN.
- The PIM-SM snooping feature assumes that the group source and the device are in different subnets and communicate through a router. The source must be in a different IP subnet than the receivers. A PIM-SM router sends PIM join and prune messages on behalf of a multicast group receiver only when the router and the source are in different subnets. When the receiver and source are in the same subnet, they do not need the router in order to find one another. They find one another directly within the subnet.

The device forwards all IPv6 multicast traffic by default. Once you enable IPv6 Multicast Traffic Reduction and PIM-SM traffic snooping, the device initially blocks all PIM-SM traffic instead of forwarding it. The device forwards PIM-SM traffic to a receiver only when the device receives a join message from the receiver. Consequently, if the source and the downstream router are in the same subnet, and PIM-SM traffic snooping is enabled, the device blocks the PIM-SM traffic and never starts forwarding the traffic. This is because the device never receives a join message from the downstream router for the group. The downstream router and group find each other without a join message because they are in the same subnet.

NOTE

If the "route-only" feature is enabled, PIM-SM traffic snooping is not supported.

Globally enabling IPv6 PIM SM traffic snooping

This feature is similar to PIM-SM traffic snooping but listens only for MLD snooping information, not PIM-SM information. You must enable both IPv6 Multicast Traffic Reduction and IPv6 PIM-SM traffic snooping to enable the device to listen for PIM-SM join and prune messages.

To enable IPv6 PIM-SM traffic snooping, enter the following commands at the global CONFIG level of the CLI.

```
device(config)# ipv6 multicast active
device(config)# ipv6 multicast pimsm-snooping
```

Syntax: `[no] ipv6 multicast [active | passive]`

When you enable IP multicast on device, all ports on the device are configured for IGMP.

If you are using *active* MLD, all ports can send MLD queries and receive MLD reports. If you are using *passive* MLD, all ports can receive MLD queries.

IPv6 Multicast Traffic Reduction cannot be disabled on individual ports of a device. IPv6 Multicast Traffic Reduction can be disabled globally by entering the `no ipv6 multicast` command.

Syntax: `[no] ipv6 multicast pimsm-snooping`

Use the *no* form of the command to disable IPv6 PIM-SM traffic snooping.

Globally configuring the query interval

If IPv6 Multicast Traffic Reduction is set to active mode, you can configure the query interval, which specifies how often a device is enabled for active IPv6 Multicast Traffic Reduction sends Group Membership queries.

NOTE

The query interval applies only to the active mode of IPv6 Multicast Traffic Reduction.

To configure the query interval, enter the following command.

```
device(config)# ipv6 multicast query-interval 120
```

Syntax: `[no] ipv6 multicast query-interval interval`

The *interval* parameter specifies the interval between queries. You can specify a value from 10 through 600 seconds. The default is 125 seconds.

Setting the MLD version

You can use the `ipv6 multicast version` command to set the MLD version (1 or 2) globally for IPv6 multicast. You can select the version of MLD by entering the following command.

```
device(config)# ipv6 multicast version 2
```

Syntax: `ipv6 multicast version version-number`

Enter 1 or 2 for the *version-number*. The default is version 1.

Configuring IPv6 multicast tracking and fast-leave

When IPv6 multicast tracking is configured, MLD fast-leave is enabled. MLD Leave Processing for MLDv1 is initiated by a receiver sending a Leave message. The router sends out a group-specific query to solicit reports from other receivers. The multicast group entry is maintained for 3 seconds to allow the processing of reports from any receivers on the LAN segment. If there are no receivers, the multicast stream is pruned.

MLD fast-leave allows a receiver to move from one multicast group to another instantly if it is the only receiver on the segment subscribed to the group. When a layer 3 switch receives a Leave message, it sends a group-specific query to see if any other receivers are present. The multicast group state is maintained for 3 seconds to process any MLD group reports from other receivers. If there are no other receivers, the multicast group entry prunes the receiver LAN segment, stopping traffic instantly.

MLDv2 also supports fast-leave processing. When an MLDv2 receiver changes the mode to Exclude, and there are no other receivers on that interface, the multicast group entry prunes the receiver LAN segment.

To enable IPv6 multicast tracking globally, enter the following command.

```
device(config)# ipv6 multicast tracking
```

Syntax: `[no] ipv6 multicast tracking`

The *no* form of this command disables the tracking process globally.

Configuring IPv6 MLD snooping on a per-VLAN basis

The following IPv6 MLD snooping parameters can be configured on a per-VLAN basis:

- Active and passive - [Configuring IPv6 multicast snooping per VLAN](#) on page 38

- MLD proxy - [Configuring MLD proxy per VLAN](#) on page 38
- PIM-SM traffic snooping - [Configuring the PIM-SM traffic snooping per VLAN](#) on page 38
- PIM proxy - [Configuring PIM proxy per VLAN](#) on page 39
- Static-group uplink - [Configuring an IPv6 multicast static group uplink per VLAN](#) on page 39
- Tracking - [Configuring IPv6 multicast tracking and fast-leave per VLAN](#) on page 40

Configuring IPv6 multicast snooping per VLAN

The **multicast6** command allows you to configure IPv6 multicast snooping parameters per VLAN. To configure IPv6 multicast snooping to VLAN 2, enter the following commands as shown in the example below.

```
device(config)# vlan 2
device(config-vlan-2)# multicast6 active
```

To remove multicast traffic reduction configurations in VLAN 2, and take the global multicast traffic reduction configuration, enter the following command. Do not enter the *active* or *passive* keywords when removing the multicast traffic reduction configuration.

```
device(config)# vlan 2
device(config-vlan-2)# no multicast6
```

Syntax: **[no] multicast6 active | passive**

When you enable IPv6 multicast for a specific VLAN, MLD snooping is enabled. The device uses MLD to maintain a table of the Group Membership reports received by the device for the specified VLAN. You can use active or passive MLD mode. There is no default mode.

The description for the MLD modes is as follows:

- **Active** - When active MLD mode is enabled, the router actively sends out MLD queries to identify IPv6 multicast groups within the VLAN and makes entries in the MLD table based on the Group Membership reports received from the network.
- **Passive** - When passive MLD mode is enabled, the router listens for MLD Group Membership reports on the VLAN specified but does not send MLD queries. The passive mode is called "MLD snooping". Use this mode when another device in the VLAN is actively sending queries.

Configuring MLD proxy per VLAN

The **multicast6 mld-proxy-enable** command enables MLD proxy for IPv6. Using the MLD proxy function, the host is able send out MLD reports on behalf of the hosts behind the switch.

To configure a device to function as an MLD proxy on VLAN 2, enter the following commands as shown in this example.

```
device(config)# vlan 2
device(config-vlan-2)# multicast6 active
device(config-vlan-2)# multicast6 mld-proxy-enable
```

Syntax: **[no] multicast6 mld-proxy-enable**

The *no* form of this command disables MLD proxy on a per-VLAN basis.

Configuring the PIM-SM traffic snooping per VLAN

In the following example, multicast traffic reduction is applied using PIM-SM traffic snooping to VLAN 2.

```
device(config)# vlan 2
device(config-vlan-2)# multicast6 active
device(config-vlan-2)# multicast6 pimsm-snooping
```

Syntax: **[no] multicast6 pimsm-snooping**

The *no* form of this command disables PIM SM traffic snooping on a per-VLAN basis.

Configuring PIM proxy per VLAN

Using the PIM proxy function, multicast traffic can be reduced by configuring a device to issue PIM join and prune messages on behalf of hosts that the configured router discovers through standard PIM interfaces. The router is then able to act as a proxy for the discovered hosts and perform PIM tasks upstream of the discovered hosts. Where there are multiple PIM downstream routers, this removes the need to send multiple messages.

When configuring PIM proxy on a VLAN, you must first configure PIM-SM traffic snooping. To configure a device to function as a PIM proxy on VLAN 2, use the following commands.

```
device(config)# vlan 2
device(config-vlan-2)# multicast6 active
device(config-vlan-2)# multicast6 pimsm-snooping
device(config-vlan-2)# multicast6 pim-proxy-enable
```

Syntax: [no] **multicast6 pim-proxy-enable**

The *no* form of this command disables PIM proxy on a per-VLAN basis.

Configuring an IPv6 multicast static group uplink per VLAN

When the **multicast6 static-group uplink** command is enabled on a snooping VLAN, the snooping device behaves like an MLD host on ports connected to the multicast router. The snooping device will respond to MLD queries from the uplink multicast PIM router for the groups and sources configured. Upon the multicast router receiving the MLD join message, it will initiate the PIM join on its upstream path towards the source to pull the source traffic down. The source traffic will stop at the MLD snooping device. The traffic will then be forwarded to the multicast receiver and router ports or dropped in hardware if no other multicast receiver and routers are present in the VLAN.

The **multicast6 static-group uplink** command can be configured under the VLAN configuration only.

The **multicast6 static-group uplink** command must be used with the **multicast6 static-group** command in order to connect a remote multicast source with the snooping VLAN where the static group is configured.

When using MLDv2, you can use the **multicast6 static-group include** or **multicast6 static-group exclude** command to statically *include* or *exclude* multicast traffic, respectively for hosts that cannot signal group membership dynamically.

To configure the snooping device to statically join a multicast group on the uplink interface, enter the following commands.

```
device(config)# vlan 10
device(config-vlan-10)# multicast6 static-group active
device(config-vlan-10)# multicast6 static-group ff2e::1 uplink
```

NOTE

The following error message will display if both static uplink MLDv1 and MLDv2 are configured for the same group: Error: Static v1 and v2 uplink configuration cannot co-exist for the same group.

To configure the physical interface Ethernet 1/1 to statically join a multicast group with an IPv6 group address of ff0e::1, enter commands such as the following.

```
device(config)# vlan 10
device(config-vlan-10)# multicast6 static-group ff0e::1 ethernet 1/1
```

To configure the snooping device to statically join a multicast stream on the uplink interface with the source address of 2003::1 in the MLDv2 include mode, enter commands such as the following.

```
device(config)# vlan 10
device(config-vlan-10)# multicast6 static-group ff1e::1 include 2003::1 uplink
```

To configure the snooping device to statically join all multicast streams on the uplink interface excluding the stream with source address 2002::1, enter commands such as the following.

```
device(config)# vlan 10
device(config-vlan-10)# multicast6 static-group active
device(config-vlan-10)# multicast6 static-group fflle::1 exclude 2002::1 uplink
```

Syntax: [no] **multicast6 static-group** *group-address* **uplink**

Syntax: [no] **multicast6 static-group** *group-address* [**include** | **exclude** *source-address*] **uplink**

The **group-address** variable specifies the group IPv6 multicast address.

The **include** or **exclude** keyword indicates a filtering action. You can specify which source (for a group) to include or exclude. The **include** or **exclude** keyword is only supported on MLDv2.

The **source-address** parameter specifies the IPv6 address of the multicast source. Each address must be added or deleted one line per source.

The **uplink** parameter specifies the port as an uplink port that can receive multicast data for the configured multicast groups. Upstream traffic will be sent to the router and will not use a port.

The **no** form of this command removes the static multicast definition. Each configuration must be deleted separately.

Configuring IPv6 multicast tracking and fast-leave per VLAN

The **multicast6 tracking** command enables IPv6 multicast tracking per VLAN. The **multicast6 tracking** command is similar to the **ipv6 multicast tracking** command. When the **multicast6 tracking** command is configured, MLD fast-leave is enabled. For more information on MLD fast-leave for IPv6 multicast tracking, refer to [Configuring IPv6 multicast tracking and fast-leave](#) on page 37.

To enable IPv6 multicast tracking per VLAN, enter commands such as the following.

```
device(config)# vlan 2
device(config-vlan-2)# multicast6 active
device(config-vlan-2)# multicast6 tracking
```

Syntax: [no] **multicast6 tracking**

The **no** form of this command disables the tracking process.

Displaying IPv6 multicast information

To display information for IPv6 multicast traffic reduction configuration, enter the following command.

```
device(config)# show ipv6 multicast
Global Multicast Traffic Reduction Configuration
MLD Snooping State : Disabled Version : 1
Group Interval : 260 Query Interval : 125
Max Response Time : 10 Robustness Var : 1
Last Member Qry Int: 5 Last Member Qry Count: 3
Querier Exp Tm : 255
MLD Proxy : Disabled Proxy Interval : 60
Filter : Disabled Tracking : Disabled
PIM Snooping : Disabled
PIM Prune Wait Time: 3
PIM Proxy : Disabled Proxy Interval : 60
VLAN snooping configurations:
VLAN ID 2
IPv6 Multicast snooping is enabled - Active. Entries 0
IPv6 Multicast MLD tracking is disabled
VLAN ID 10
IPv6 Multicast snooping is enabled - Active. Entries 0
IPv6 Multicast pimsm snooping is enabled
```


Syntax: `show ipv6 multicast [mldv2 vlan-id | pim vlan-id | static vlan-id | statistics vlan-id | tracking vlan-id | vlan vlan-id]`

The **mldv2** *vlan-id* parameter displays IPv6 multicast information specific to MLDv2 for a specified port-based VLAN.

The **pim** *vlan-id* parameter displays IPv6 multicast information specific to PIM for a specified port-based VLAN.

The **static** *vlan-id* parameter displays information for the number of static MLD snooping entries configured in a specified port-based VLAN.

The **statistics** *vlan-id* parameter displays IPv6 multicast statistics for a specified port-based VLAN.

The **tracking** *vlan-id* parameter displays tracking information for MLDv2 hosts for a specified port-based VLAN.

The **vlan** *vlan-id* parameter displays IPv6 multicast PIM information for a specified port-based VLAN. The *vlan-id* variable is entered in decimal format.

Table 3 describes the fields displayed by the **show ipv6 multicast** command.

TABLE 3 Output from the **show ipv6 multicast** command

Field	Description
Global Multicast Traffic Reduction Configuration	Indicates all multicast traffic configuration displayed for all parameters.
MLD Snooping State:	Indicates whether MLD snooping is enabled or disabled. If enabled, it indicates if the feature is configured as passive or active.
Global Interval	Indicates the time until the groups time out if no reports are received.
Max Response Time	The length of time in seconds that the router will wait for an IGMP (V1 or V2) response from an interface before concluding that the group member on that interface is down and removing it from the group.
Last Member Qry Int	Indicates when a leave is received; a group-specific query is sent. The last member query count is the number of queries with a time interval of (LMQT) is sent.
Querier Exp Tm	Indicates the time until the querier times out if no query is received.
MLD Proxy	Indicates whether MLD Proxy is enabled or disabled. If enabled, it indicates if the feature is configured as passive or active.
Filter	Indicates whether filtering is enabled or disabled. Filtering multicast groups allows the device to filter out all multicast traffic groups other than the ones for which the device has received Group Membership reports.
PIM Snooping	Indicates if PIM snooping is enabled. If disabled, this line does not appear.
PIM Prune Wait Time	The amount of time a PIM router will wait before stopping traffic to neighbor routers that do not want the traffic. The value can be from 0 to 3 seconds. The default is 3 seconds.
PIM Proxy	Indicates whether PIM proxy is enabled or disabled. If enabled, it indicates if the feature is configured as passive or active.
Version	The MLD version (1 or 2) operating on the router.
Query Interval	How often the router will query an interface for group membership.
Robustness Var	Used to fine tune for unexpected loss on the subnet. The value is used to calculate the group interval.
Last Member Qry Count	Specifies the number of group-specific queries when a leave is received.
Proxy Interval	Indicates the time interval in seconds between PIM proxy and join messages. You can specify a value from 10 through 600 seconds.
Tracking	Indicates whether tracking is enabled or disabled.
VLAN ID	The port-based VLAN to which the information listed below the ID applies.
IPv6 multicast traffic snooping	Indicates whether IPv6 multicast traffic snooping is enabled or disabled. If enabled, it indicates if the feature is configured as passive or active.

TABLE 3 Output from the **show ipv6 multicast** command (continued)

Field	Description
IPv6 Multicast MLD tracking	Indicates whether IPv6 multicast MLD tracking is enabled or disabled.
IPV6 Multicast pimsm snooping	Indicates whether pimsm snooping is enabled or disabled.

To display detailed IPv6 multicast traffic reduction information for a specified VLAN, enter the following command at any level of the CLI.

```
device# show ipv6 multicast vlan 180
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
VLAN State Mode      Active    Time      (*, G) (S, G)
      Querier Query Count Count
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
180  Ena   Active    Self 124    8        8
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Router ports:
Flags: R-Router Port, V1|V2: MLD Receiver, P_G|P_SG: PIM Join
1   (*, ffle::6 ) 00:07:48 NumOIF: 1 profile: 8
   Outgoing Interfaces:
     e4/12 vlan 180 ( V1 ) 00:07:43/126s
1   (2001:1:1:1::180:101, ffile::6) in e4/9 vlan 180 00:07:48 NumOIF:1 profile: 8
   Outgoing Interfaces:
     e4/12 vlan 180 ( V1 ) 00:07:43/0s
     FID: 0x8015 MVID: None
2   (*, ffile::4 ) 00:07:48 NumOIF: 1 profile: 8
   Outgoing Interfaces:
     e4/12 vlan 180 ( V1 ) 00:07:43/126s
```

Syntax: **show ipv6 multicast** [**vlan** *vlan-id*]

The **vlan** *vlan-id* parameter displays IPv6 multicast PIM information for a specified port-based VLAN. The *vlan-id* variable is entered in decimal format.

Table 4 describes the output parameters of the **show ipv6 multicast vlan** command.

TABLE 4 Output parameters of the **show ipv6 multicast vlan** command

Field	Description
VLAN	Shows the ID of the configured VLAN.
State	Shows whether the VLAN interface is enabled or disabled.
Mode	Shows whether the VLAN interface is in active mode or passive mode.
Active Querier	Shows the active IGMP querier for the VLAN.
Time Query	Shows the time countdown to generate the next query message.
(*, G)Count	Shows the count of (*,G) entries.
(S, G)Count	Shows the count of (S,G) entries.
Flags	Shows the interface flag for the entry.
V1 V2	Shows the version of the IGMP message received.
P_G	Indicates that a PIM (*,G) join was received on that interface.
P_SG	Indicates that a PIM (S,G) join was received on that interface.
NumOIF	Show the count of the outgoing interfaces.
profile	Shows the profile ID associated with the stream.
Outgoing Interfaces	Shows the outgoing interfaces.
FID	Shows the FID resource allocated for a particular entry.
MVID	Shows the MVID resource allocated for a particular entry.

IPv4 Multicast Routing

• Overview of IP multicasting.....	43
• Changing global IP multicast parameters.....	44
• Mtrace overview.....	47
• Support for Multicast Multi-VRF.....	49
• Adding an interface to a multicast group.....	50
• Multicast non-stop routing.....	51
• Passive Multicast Route Insertion (PMRI)	53
• IP multicast boundaries.....	54
• Configuring Layer 3 Multicast filter for the hardware.....	62
• PIM Dense	65
• PIM Sparse	80
• Multicast Outgoing Interface (OIF) list optimization.....	91
• Configuring Multicast Source Discovery Protocol (MSDP).....	107
• Configuring MSDP mesh groups	120
• MSDP Anycast RP.....	122
• PIM Anycast RP.....	126
• PIM over MCT intermediate router functionality.....	128
• Configuring a static multicast route.....	139
• IGMP V3.....	141

Overview of IP multicasting

Multicast protocols allow a group or channel to be accessed over different networks by multiple stations (clients) for the receipt and transmission of multicast data.

Distribution of stock quotes, video transmissions such as news services and remote classrooms, and video conferencing are all examples of applications that use multicast routing.

Extreme devices support Protocol-Independent Multicast (PIM) protocol, along with the Internet Group Membership Protocol (IGMP).

PIM is broadcast and pruning multicast protocol that deliver IP multicast datagrams. This protocol employs reverse path lookup check and pruning to allow source-specific multicast delivery trees to reach all group members. PIM builds a different multicast tree for each source and destination host group.

PIM can concurrently operate on different ports of a device. The CAM can hold up to 1535 IPv4 multicast entries.

Multicast terms

The following terms are commonly used in discussing multicast-capable devices. These terms are used throughout this chapter:

Node: Refers to a device.

Root Node: The node that initiates the tree building process. It is also the device that sends the multicast packets down the multicast delivery tree.

Upstream: Represents the direction from which a device receives multicast data packets. An upstream device is a node that sends multicast packets.

Downstream : Represents the direction to which a device forwards multicast data packets. A **downstream** device is a node that receives multicast packets from upstream transmissions.

Group Presence : Means that a multicast group has been learned from one of the directly connected interfaces. Members of the multicast group are present on the device.

Intermediate nodes : Devices that are in the path between source devices and leaf devices.

Leaf nodes: Devices that do not have any downstream devices.

Multicast Tree: A unique tree is built for each source group (S,G) pair. A multicast tree is comprised of a root node and one or more nodes that are leaf or intermediate nodes.

Changing global IP multicast parameters

The following sections apply to PIM-DM, PIM-SM and IGMP.

Concurrent support for multicast routing and snooping

Multicast routing and multicast snooping instances work concurrently on the same device. For example, you can configure PIM routing on certain VEs interfaces and snooping on other VEs or VLANs. The limitation is that either multicast snooping or routing can be enabled on a VE interface or VLAN, but not on both. This is because all of the multicast data and control packets (IGMP, PIM) received on the snooping VLAN are handled by multicast snooping and do not reach the multicast routing component. Similarly, any multicast data or control packets received on a VE interface enabled with PIM routing are handled by the PIM, routing component and are not seen by the IGMP or PIM snooping component.

The following considerations apply when configuring concurrent operation of Multicast Routing and Snooping.

1. Either multicast snooping or routing can be enabled on a VE or VLAN but not both.
2. There may be slight multicast traffic loss on one receiver while configuring **ip multicast no-fid-updates**, when other receiver sends IGMP leave.
3. Snooping can be enabled globally (**ip multicast active / passive**) as well **multicast routing (ip multicast-routing)** .
4. The global snooping configuration is inherited by all current VLANs that are not enabled for multicast routing.
5. The global snooping configuration is also inherited by all new VLANs. Enabling multicast routing on a newly created VLAN or VE automatically disables snooping on the VLAN or VE.
6. When a VLAN-level snooping is configured, it is displayed.

Defining the maximum number of PIM cache entries

You can use the following run-time command to define the maximum number of repeated PIM traffic being sent from the same source address and being received by the same destination address. To define this maximum for the default VRF, enter the following commands.

```
device(config)# router pim
device(config-pim-router)# max-mcache 999
```

Syntax: [no] **max-mcache** *num*

The *num* variable specifies the maximum number of multicast cache entries for PIM in the default VRF. If not defined by this command, the maximum value is determined by available system resources.

To define the maximum number of PIM Cache entries for a specified VRF, use the following command.

```
device(config)# router pim vrf vpn1
device(config-pim-router-vrf-vpn1)# max-mcache 999
```

Syntax: `[no] router pim [vrf vrf-name]`

Syntax: `[no] max-mcache num`

The **vrf** parameter specified with the **router pim** command allows you to configure the **max-mcache** command for a virtual routing instance (VRF) specified by the variable *vrf-name*.

The *num* variable specifies the maximum number of multicast cache entries for PIM in the specified VRF. If not defined by this command, the maximum value is determined by available system resources.

Defining the maximum number of multicast VRF CAM entries

To use a **run time** command to set the maximum values for multicast VRF CAM entries for all VRFs or for a specified VRF.

Defining the maximum number of multicast VRF CAM entries for all VRFs

You can use the following **run-time** command to define the maximum number of multicast VRF CAM entries by entering a command such as the following.

```
device(config)# ip multicast-max-all-vrf-cam 3072
```

Syntax: `[no] ip multicast-max-all-vrf-cam num`

The *num* variable specifies the maximum number of multicast VRF CAM entries for all VRFs. This setting does not effect the default VRF. The maximum possible value is 32780 and the default value is 2048.

Defining the maximum number of multicast VRF CAM entries for a specified VRF

You can use the following run-time command to define the maximum number of multicast VRF CAM entries for a specified VRF by entering commands such as the following.

```
device(config)# vrf vpn1
device(config-vrf-vpn1)# ip multicast-max-cam 3072
```

Syntax: `[no] vrf vrf-name`

Syntax: `[no] ip multicast-max-cam num`

The **vrf** parameter specifies the virtual routing instance (VRF) specified by the variable *vrf-name*.

The *num* variable specifies the maximum number of multicast VRF CAM entries for the specified VRF. This setting can be any number up to the limit set using the **ip multicast-max-all-vrf-cam** command.

Defining the maximum number of IGMP group addresses

You can use the following **run-time** command to set the maximum number of IGMP addresses for the default VRF. To define this maximum for the default VRF, enter the following command.

```
device(config)# ip igmp max-group-address 1000
```

Syntax: `[no] ip igmp max-group-address num`

The *num* variable specifies the maximum number of IGMP group addresses you want to make available for the default VRF. If not defined by this command, the maximum value is determined by available system resources.

Changing IGMP V1 and V2 parameters

IGMP allows Extreme devices to limit the multicast of IGMP packets to only those ports on the device that are identified as IP Multicast members.

The device actively sends out host queries to identify IP Multicast groups on the network, inserts the group information in an IGMP packet, and forwards the packet to IP Multicast neighbors.

The following IGMP V1 and V2 parameters apply to PIM:

- **IGMP query interval** - Specifies how often the Extreme device queries an interface for group membership. Possible values are 2 - 3600. The default is 125.
- **IGMP group membership time** - Specifies how many seconds an IP Multicast group can remain on a Extreme device interface in the absence of a group report. Possible values are 5 - 26000. The default is 260.
- **IGMP maximum response time** - Specifies how many seconds the Extreme device will wait for an IGMP response from an interface before concluding that the group member on that interface is down and removing the interface from the group. Possible values are 1 - 25. The default is 10.

To change these parameters, you must first enable IP multicast routing by entering the following CLI command at the global CLI level.

```
device(config)# ip multicast-routing
```

Syntax: **[no] ip multicast-routing**

NOTE

You must enter the **ip multicast-routing** command before changing the global IP Multicast parameters. Otherwise, the changes do not take effect and the software uses the default values. Also, entering **no ip multicast-routing** will reset all parameters to their default values.

Modifying IGMP (V1 and V2) query interval period

The IGMP query interval period defines how often a device will query an interface for group membership. Possible values are 2 - 3600 seconds and the default value is 125 seconds.

To modify the default value for the IGMP (V1 and V2) query interval, enter the following.

```
device(config)# ip igmp query-interval 120
```

Syntax: **[no] ip igmp query-interval num**

The *num* variable specifies the number of seconds and can be a value from 2 - 3600.

The default value is 125.

Modifying IGMP (V1 and V2) membership time

Group membership time defines how long a group will remain active on an interface in the absence of a group report. Possible values are from 5 - 26000 seconds and the default value is 260 seconds.

To define an IGMP (V1 and V2) membership time of 240 seconds, enter the following.

```
device(config)# ip igmp group-membership-time 240
```

Syntax: **[no] ip igmp group-membership-time num**

The *num* variable specifies the number of seconds and can be a value from 5 - 26000.

The default value is 260.

Modifying IGMP (V1 and V2) maximum response time

Maximum response time defines how long the Extreme device will wait for an IGMP (V1 and V2) response from an interface before concluding that the group member on that interface is down and removing the interface from the group. Possible values are 1 - 25. The default is 10.

To change the IGMP (V1 and V2) maximum response time, enter a command such as the following at the global CONFIG level of the CLI.

```
device(config)# ip igmp max-response-time 8
```

Syntax: [no] ip igmp max-response-time *num*

The *num* variable specifies the number of seconds and can be a value from 1 - 25. The default is 10.

Security Enhancement for IGMP

A security enhancement has been made to IGMPv2 to adhere to the following recommendation of RFC 2236: "Ignore the Report if you cannot identify the source address of the packet as belonging to a subnet assigned to the interface on which the packet was received."

NOTE

When used in applications such as IP-TV (or any multicast application in general), the administrator should ensure that the set-top box (or multicast client) is configured on the same subnet as the v.e. configured on the device. This is typically the case but is emphasized here to ensure correct operation. Without this configuration, IGMP messages received by the device are ignored which causes an interruption in any multicast traffic directed towards the set-top box (multicast client).

Mtrace overview


"mtrace" is a diagnostic tool to trace the multicast path from a specified source to a destination for a multicast group. It runs over IGMP protocol. Mtrace uses any information available to it to determine a previous hop to forward the trace towards the source.

There are three main components in an mtrace implementation. They are mtrace query, mtrace request, and mtrace response.

The unicast "traceroute" program allows the tracing of a path from one machine to another. The key mechanism for unicast traceroute is the ICMP TTL exceeded message, which is specifically excluded as a response to multicast packets. The multicast traceroute facility allows the tracing of an IP multicast routing path. Multicast traceroute also requires special implementations on the part of routers.

Multicast traceroute uses any information available to it in the router to determine a previous hop to forward the trace towards the source. Multicast routing protocols vary in the type and amount of state they keep; multicast traceroute endeavors to work with all of them by using whatever is available. For example, if a PIM-SM router is on the (*,G) tree, it chooses the parent towards the RP as the previous hop. In these cases, no source/group-specific state is available, but the path may still be traced.

FIGURE 6 Network topology



Mtrace components

There are 3 main components in a multicast traceroute implementation. They are:

1. Mtrace Query
 2. Mtrace Request
 3. Mtrace Response
- Mtrace Query

The party requesting the traceroute sends a traceroute query packet to the last-hop multicast router for the given destination. The query and request have the same opcode, the receiving router can distinguish between a query and a request by checking the size of the packet. A query is a request packet with none of the response fields filled up.

- Mtrace Request

The last-hop router turns the Query packet into a Request packet by adding a response data block containing its interface addresses and packet statistics, and then forwards the Request packet via unicast to the router that it believes is the proper previous hop for the given source and group. Each hop adds its response data to the end of the Request packet, then unicast forwards it to the previous hop.

- Mtrace Response

The first hop router (the router that believes that packets from the source originate on one of its directly connected networks) changes the packet type to indicate a Response packet and sends the completed response to the response destination address. The response may be returned before reaching the first hop router if a fatal error condition such as "no route" is encountered along the path.

Configuring mtrace

To explain how mtrace works, let's take the network topology depicted in [Mtrace overview](#) on page 47. The mtrace can be started on any router on the network. The format of the command to start a mtrace would be:

```
device#mtrace ipv6 source 102::1 destination 101::2 group ff1d::2
Mtrace handle query from src 102::1 to dest 101::2 through group ff1d::2
```

Collecting Statistics, waiting time 5 seconds.....

```
Type Control-c to abort
0 12::1 PIM thresh^ 1 MTRACE_NO_ERR
1 13::1 PIM thresh^ 1 MTRACE_NO_ERR
2 102::2 PIM thresh^ 1 MTRACE_REACHED_RP
```

Syntax: `mtrace [ipv6] [vrf] [vrf name] source ip-address [destination ip-address] [group ip-address]`

TABLE 5 parameters of the mtrace command

Field	Description
Source	IP address of the Multicast capable source. This is a unicast address of the beginning of the path to be traced.
Destination	Address of the unicast destination. If omitted, the trace starts from the system where the command was issued.
Group	Multicast address of the group to be traced. Default address is 224.2.0.1 for IPv4 and FF0E:0:0:0:0:0:10E (IETF-2_AUDIO) for IPv6.

Assume that the destination is 195.1.2.1, source is 192.168.1.1 and group is 225.1.1.1.

The mtrace query is initially sent from R7. The initial header is not to be modified by any of the routers. R5 adds a response block based on the (S, G) or the (*, G) entry and adds its incoming interface, outgoing interface and other information specified in the draft and sends it to its upstream neighbor which is R6. R6 similarly adds a response block and sends it to its upstream neighbor R2, likewise till it reaches R1. Once it reaches R1, R1 determines that it is the first hop router and completes the response block and sends the response back to R7. R7 now reads the information from the packet and prints it out.

Support for Multicast Multi-VRF

Multicast Multi-VRF support for the Extreme device includes the following:

- **Static Mroute** - As described in [Configuring a static multicast route within a VRF](#) on page 140 you can configure static multicast route from within a specified VRF.
- **PIM (PIM-SM and PIM-DM)** - The procedure for configuring PIM within a VRF instance is described in [Enabling PIM for a specified VRF](#) on page 69 and [Enabling PIM Sparse for a specified VRF](#) on page 83.

System max parameter changes

Several changes to the **system max** commands have been made in support of Multicast Multi-VRF. That includes retiring the following system max commands:

system-max multicast-route

system-max pim-mcache

system-max igmp-max-group-address

These commands which require a system reload to take effect have been replaced by the following runtime commands:

ip max-mroute - This command replaces the **system-max multicast-route** command.

max-mcache - This command described in [Defining the maximum number of PIM cache entries](#) on page 44 replaces the **system-max pim-mcache** command.

ip igmp max-group-address - This command described in [Defining the maximum number of IGMP group addresses](#) on page 45 replaces the **system-max igmp-max-group-address** command.

NOTE

If the deprecated **system-max** commands are used, the new runtime commands will be substituted in the running config.

Additionally, you can set a maximum value for multicast VRF CAM entries as described in [Defining the maximum number of multicast VRF CAM entries](#) on page 45.

Show and clear command support

The following **show** and **clear** commands have been introduced or enhanced to support Multicast Multi-VRF:

- clear ip igmp cache
- clear ip igmp traffic
- show ip igmp group
- show ip igmp interface
- show ip igmp settings
- show ip igmp traffic

Adding an interface to a multicast group

You can manually add an interface to a multicast group. This is useful in the following cases:

- Hosts attached to the interface are unable to add themselves as members of the group using IGMP.
- There are no members for the group attached to the interface.

When you manually add an interface to a multicast group, the device forwards multicast packets for the group but does not itself accept packets for the group.

You can manually add a multicast group to individual ports only. If the port is a member of a virtual routing interface, you must add the ports to the group individually.

To manually add a port to a multicast group, enter a command such as the following at the configuration level for the port.

```
device(config-if-e10000-1/1)# ip igmp static-group 224.2.2.2
```

This command adds port 1/1 to multicast group 224.2.2.2.

To add a port that is a member of a virtual routing interface to a multicast group, enter a command such as the following at the configuration level for the virtual routing interface.

```
device(config-vif-1)# ip igmp static-group 224.2.2.2 ethernet 5/2
```

This command adds port 5/2 in virtual routing interface 1 to multicast group 224.2.2.2.

Syntax: [no] ip igmp static-group *ip-addr* [ethernet *slot/portnum*]

The `ip-addr` parameter specifies the group number.

The `ethernet slot/portnum` parameter specifies the port number. Use this parameter if the port is a member of a virtual routing interface, and you are entering this command at the configuration level for the virtual routing interface.

Manually added groups are included in the group information displayed by the following commands:

- `show ip igmp group`
- `show ip pim group`

Multicast non-stop routing

Multicast non-stop routing (NSR) provides hitless upgrade and switchover support for all IPv4 multicast, including default and non-default VRFs for IPv4 PIM-DM, PIM-SM, and PIM-SSM. Multicast NSR is not supported for IPv6 multicast. The software multicast state is kept in sync between the active and standby MPs. As the Extreme system enters a hitless upgrade or switchover state, the standby MP will take over as the new active MP. The new active MP will carry a pre-installed multicast state that was originally supported by the previous MP. The new active MP will revalidate the pre-installed multicast state, and pick up any new changes as needed before marking the multicast state as operational. When the LP is ready to complete the hitless upgrade or switchover process, the operational multicast state will be downloaded to the LP CPU. When the LP resets, and the outage of the LP CPU occurs, pre-existing hardware forwarding multicast traffic will continue to flow without disruption, and the hardware multicast forwarding state is retained in the LP hardware.

Multicast NSR is globally enabled across all VRFs by configuring the `ip multicast-nonstop-routing` command. For more information on configuring the `ip multicast-nonstop-routing` command, refer to [Configuring multicast non-stop routing](#) on page 51.

NOTE

During hitless reload, if any changes occur to the existing multicast forwarding records, then multicast receivers of the same forwarding records may see traffic loss.

NOTE

Hitless upgrade support for multicast NSR is supported only on XMR Series and MLX Series devices.

Configuration considerations

- Multicast NSR is not supported for IPv6 multicast and layer 2 multicast.
- When multicast NSR is turned on, unicast routing must be protected by NSR or graceful restart on all multicast VRFs.
- Any multicast flow that does not have a hardware CAM entry programmed prior to hitless upgrade or switchover will not be protected under multicast NSR. The multicast entry for such a flow shall be recreated upon the completion of the NSR process.

Configuring multicast non-stop routing

To globally enable multicast non-stop routing for all VRFs, enter the `ip multicast-nonstop-routing` command on the CLI as shown in the example below.

```
device(config)#ip multicast-nonstop-routing
```

Syntax: ip multicast-nonstop-routing

During a hitless upgrade and switchover on the MP, the following syslog message is generated on the CLI.

```
Feb 3 14:09:58 Mcastv4 detected MP switchover, set switchover in progress to TRUE
Feb 3 14:10:07 Mcastv4 confirms unicast RTM is ready
Feb 3 14:10:07 Mcastv4 switchover done, set switchover in progress mode to FALSE
```

The syslog message displayed above shows the state transition of multicast NSR as the standby MP takes over as the active MP. The multicast data traffic will continue to flow during state transition.

Displaying the multicast NSR status

To display the multicast NSR status, enter the following command.

```
device#show ip pim nsr
Global Mcast NSR Status
NSR: ON
Switchover In Progress Mode: FALSE
Dy-Sync Postpone Flag: FALSE
```

The following table displays the output from the **show ip pim nsr** command.

TABLE 6 Output from the **show ip pim nsr**

This field...	Displays...
NSR	The NSR field indicates if the ip multicast-nonstop-routing command is enabled (ON) or disabled (OFF).
Switchover in Progress Mode	The Switchover in Progress Mode field indicates if the multicast traffic is in the middle of a switchover (displaying a TRUE status), or not (displaying a FALSE status).
Dy-Sync Postpone Flag	After the current switchover or hitless upgrade is complete, an update to the batched dy-sync may or may not need posting.

Displaying counter and statistic information for multicast NSR

To display multicast NSR counter and statistics information from the MP, enter the following command.

```
device# show ip pim counter nsr
Mcache sync (entity id: 203)
  pack: 0
  unpack: 0
  ack: 0
RPset sync (entity id: 201)
  pack: 0
  unpack: 0
  ack: 0
BSR status (entity id: 202)
  pack: 1
  unpack: 0
  ack: 1
```

Syntax: **show ip pim [vrf vrf_name] counter nsr**

The **vrf** parameter allows you to display IP PIM counters for the VRF instance specified by the *vrf-name* variable.

The following table displays the output from the **show ip pim counter nsr** command.

TABLE 7 Output from the **show ip pim counter nsr** command

This field...	Displays...
Mcache sync	The mcache NSR sync queue that carries the NSR sync message for mcache updates.
pack	The number of NSR sync messages that are packed from current MP to the other MP.
unpack	The number of NSR sync messages that are received and unpacked by the current MP.

TABLE 7 Output from the `show ip pim counter nsr` command (continued)

This field...	Displays...
ack	The number of NSR sync acknowledgement the current MP received.
RPset sync	The RPset sync queue that carries the NSR sync message for RPset update.
BSR status	The BSR status sync queue that carries the NSR sync message for BSR information update.

Passive Multicast Route Insertion (PMRI)

To prevent unwanted multicast traffic from being sent to the CPU, PIM Routing and Passive Multicast Route Insertion (PMRI) can be used together to ensure that multicast streams are only forwarded out ports with interested receivers and unwanted traffic is dropped in hardware on Layer 3 Switches.

PMRI enables a Layer 3 switch running PIM Sparse to create an entry for a multicast route (e.g., (S,G)), with no directly attached clients or when connected to another PIM device (transit network).

When a multicast stream has no output interfaces, the Layer 3 Switch can drop packets in hardware if the multicast traffic meets either of the following conditions:

In PIM-SM:

- The route has no OIF *and*
- If directly connected source passed source RPF check *and* completed data registration with RP *or*
- If non directly connected source passed source RPF check.

In PIM-DM:

- The route has no OIF *and*
- passed source RPF check *and*
- Device has no downstream PIM neighbor.

If the OIF is inserted after the hardware-drop entries are installed, the hardware entries will be updated to include the OIFs.

NOTE

Disabling hardware-drop does not immediately take away existing hardware-drop entries, they will go through the normal route aging processing when the traffic stops.

Configuring PMRI

PMRI is enabled by default. To disable PMRI, enter commands such as the following.

```
device(config)# router pim
device(config-pim-router) # hardware-drop-disable
```

Syntax: `[no] hardware-drop-disable`

Displaying hardware-drop

Use the `show ip pim sparse` command to display if the hardware-drop feature has been enabled or disabled.

```
device(config)#show ip pim sparse
Global PIM Sparse Mode Settings
```

```

Hello interval      : 30           Neighbor timeout          : 105
Bootstrap Msg interval: 60         Candidate-RP Advertisement interval: 60
Join/Prune interval : 60           SPT Threshold            : 1
Inactivity interval : 180          SSM Enabled              : No
Hardware Drop Enabled
: Yes
show ip pim sparse

```

IP multicast boundaries

The Multicast Boundary feature is designed to selectively allow or disallow multicast flows to configured interfaces.

NOTE

Beginning release 5.7, the IP multicast boundaries feature is available only for backward compatibility. Brocade recommends using the Layer 3 multicast filter for the hardware feature in place of the IP multicast boundaries feature.

The **ip multicast-boundary** command allows you to configure a boundary on PIM enabled interface by defining which multicast groups may not forward packets over a specified interface. This includes incoming and outgoing packets. By default, all interfaces that are enabled for multicast are eligible to participate in a multicast flow provided they meet the multicast routing protocol's criteria for participating in a flow.

Configuration considerations

The configuration considerations are as follows:

- Only one ACL can be bound to any interface.
- Normal ACL restrictions apply as to how many software ACLs can be created, but there is no hardware restrictions on ACLs with this feature.
- Creation of a static IGMP client is allowed for a group on a port that may be prevented from participation in the group on account of an ACL bound to the port's interface. In such a situation, the ACL would prevail and the port will not be added to the relevant entries.
- Either standard or extended ACLs can be used with the multicast boundary feature. When a standard ACL is used, the address specified is treated as a group address and NOT a source address.
- When a boundary is applied to an ingress interface, all packets destined to a multicast group that is filtered out will be dropped by software. Currently, there is no support to drop such packets in hardware.
- The **ip multicast-boundary** command may not stop clients from receiving multicast traffic if the filter is applied on the egress interface up-stream from RP.

Configuring multicast boundaries

Multicast boundaries can be configured for IPv4 or IPv6.

To define boundaries for PIM enabled interfaces, enter a commands such as the following.

```

device(config)# interface ve 40
device(config-vif-40)#ip multicast-boundary MyBrocadeAccessList

```

Multicast boundaries can be configured for IPv6 as shown in the following.

```

device(config)# interface ethernet 1/2
device(config-if-e1000-1/2)#ipv6 multicast-boundary MyBrocadeAccessList

```

Syntax: **[no]** ip multicast-boundary *acl-spec*

Syntax: `[no] ipv6 multicast-boundary acl-spec`

Use the *acl-spec* parameter to define the number or name identifying an access list that controls the range of group addresses affected by the boundary.

Use the `no ip multicast boundary` command to remove the boundary on a PIM enabled interface.

The ACL, MyBrocadeAccessList can be configured using standard ACL syntax. ACLs are described in Brocade NetIron Security Configuration Guide however, some examples of how ACLs can be used to filter multicast traffic are provided below:

Standard ACL to permit multicast traffic

To permit multicast traffic for group 225.1.0.2 and deny all other traffic, enter the following command.

```
device(config)# access-list 10 permit host 225.1.0.2
device(config)# access-list 10 deny any
```

Extended ACL to deny multicast traffic

To deny multicast data traffic from group 225.1.0.1 and permit all other traffic,

```
device(config)# access-list 101 deny ip any host 225.1.0.1
device(config)# access-list 101 permit ip any any
```

Extended ACL to permit multicast traffic

To permit multicast data traffic from source 97.1.1.50 for group 225.1.0.1 and deny all other traffic,

```
device(config)# access-list 102 permit ip host 97.1.1.50 host 225.1.0.1
device(config)# access-list 102 deny ip any any
```

Displaying multicast boundaries

To display multicast boundary information, use the `show ip pim interface` command.

```
device# show ip pim interface
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Interface|Local      |Mode|Ver|Designated Router |TTL|Multicast| VRF  | DR
      |Address    |   |   |Address           |Thr|Boundary |     | Prio
-----+-----+-----+-----+-----+-----+-----+-----+-----+
v10    10.1.2.1   SM  V2  Itself           1  None    default 30
v30    123.1.1.2  SM  V2  Itself           1  None
v40    124.1.1.2  SM  V2  Itself           1  101
```

Syntax: `show ip pim [vrf vrf-name] interface [ethernet slot/portnum | ve num | tunnel num]`

The **vrf** option allows you to display multicast boundary information for the VRF instance identified by the *vrf-name* variable.

The **ethernet** *port-number* parameter specifies the physical port.

The **ve** *num* parameter specifies a virtual interface.

The **tunnel** *num* parameter specifies a GRE tunnel interface that is being configured. The GRE tunnel interface is enabled under the device PIM configuration.

Performing IPv4 Multicast RPF shortcut using LSP paths

This feature provides the ability to run multicast IPv4 routing protocols (PIM-SM, SSM, PIM-DM) using native IPv4 multicast forwarding when IGP (OSPF, IS-IS) shortcut feature or IPoMPLS is enabled. When enabled, RPF (Reverse Path Forwarding) lookup results in a shortcut route (Label Switched Path or LSP). Similarly RPF lookup results in a BGP route using MPLS tunnel or a static route point to LSP end point. The LSP path resulted from these lookup cannot be directly used for RPF operations in multicast. This feature aids in RPF path resolution when the RPF lookup results in LSP.

When this feature is enabled, RPF lookup ignores the LSP route and uses the underlying native route as the RPF path. In unicast routing the LSP path is used for forwarding. In multicast routing the underlying native route is used.

FIGURE 7 LSP path topology




Figure 7 shows the RPF lookup to the source, which could result in the LSP paths between R1 and R2. This feature uses native routes R1-R3-R2 and R1-R4-R2, and ECMP logic decides the RPF shortcut.

This feature is useful when the core is running MPLS and multicast routing protocols (like GRE MVPN). There can be an IGP shortcut path or LSP path which would provide next-hop as the tunnel interface. The IGP can route unicast traffic over these tunnels to destinations that are downstream from the egress router of the tunnel. Without this feature, RPF path through LSP cannot be directly used to send PIM control packets. With this feature enabled, Multicast control the packets or traffic would use LSP's underlying native path as RPF Path. If this feature is not enabled RPF resolution would fail in these scenarios.

Limitations

These are the limitations in determining RPF shortcut.

1. No support for IPv6 multicast
2. No VRF support, only default VRF as MPLS doesn't have VRF support
3. No Support for Static Mroute to LSP tunnel end point

4. No IGP shortcut support for LDP tunnels
5. No support for ECMP between LSP path and non-LSP path
6. No NSR support for shortcut routes as there is no support in RSVP for NSR
7. No IS-IS announce metric feature support
8. No support for recursive BGP shortcut lookup

Enabling RPF shortcut feature

Use this command to enable RPF shortcut for LSP paths. If IGP shortcut was enabled prior to enabling multicast RPF shortcut, multicast cache entries whose RPF lookup resulted in shortcut would be marked with **PIM_FWD_NEED_REROUTE** as there won't be any PIM neighbor on the MPLS Interface. If RPF lookup results in the LSP path, then another lookup is done to get the underlying native route and that route's next-hop is used as the RPF.

```
device(config)# router pim
device(config-pim-router)# rpf shortcut
```

After this command is executed, if IGP shortcut is not enabled, then there won't be any LSP paths to be considered for the RPF lookup. Later, when IGP Shortcut is enabled and LSP path becomes best path, then change is handled through route change notification.

If there are ECMP LSP paths, ECMP path is chosen based on the multicast ECMP path selection logic. Enabling this command helps MBGP to accept MPLS tunnel as valid next-hop destination and to install the route into MRTM.

Use this command, when you disable multicast RPF shortcut feature. It reroutes the multicast cache entries. RPF for the multicast cache entries return either LSP or non-LSP path depends on IGP shortcut enabling. This command is only supported for default VRF.

```
device(config)# router pim
device(config-pim-router)# no rpf shortcut
```

NOTE

It is recommended to have this multicast RPF shortcut feature enabled when IGP shortcut is enabled.

If this feature is disabled, configuration will be removed from multicast and if the native route changes then the RPF for the multicast cache entries would change. Further, RPF lookup after disabling this feature will not use LSP's underlying native path as RPF path. If RPF lookup results in non-LSP multicast enabled path, RPF would resolve or else RPF resolution would fail.

Syntax: [no] rpf shortcut

Displaying the RPF routing path

This show command displays the RPF information, indicates that the route was learned through which LSP path, if the RPF happens to be through LSP path.

```
device(config)# show ip pim rpf source-address group-address
device(config)# show ip pim rpf 130.50.11.10 226.10.10.1
Upstream LSP1 nbr 55.55.55.55 on e4/1
```

Displaying flag for RPFS for multicast cache entries

mcache show command output in MP is modified to show the flag RPFS if the resolution is through LSP for the multicast cache entries using shortcut as the RPF and also, it displays a flag to indicate that this entry is using RPF shortcut (LSP path).

```
device(config)# show ip pim mcache
IP Multicast Mcache Table
Entry Flags : SM - Sparse Mode, SSM - Source Specific Multicast, DM - Dense Mode
              RPT - RPT Bit, SPT - SPT Bit, LSRC - Local Source, LRCV - Local Receiver,
              RPFS - RPF Shortcut
```


```
, HW - HW Forwarding Enabled, FAST - Resource Allocated,
    TAG - Need For Replication Entry, REGPROB - Register In Progress,
    REGSUPP - Register Suppression Timer, MSDPADV - Advertise MSDP,
    NEEDRTE - Route Required for Src/RP, PRUN - DM Prune Upstream
Interface Flags: IM - Immediate, IH - Inherited, WA - Won Assert
    MJ - Membership Join, MI - Membership Include, ME - Membership Exclude
    BR - Blocked RPT, BA - Blocked Assert, BF - Blocked Filter, BI - Blocked IIF,
    BM - Blocked MCT
Total entries in mcache: 1
(54.1.1.10, 226.0.1.0) in v52 (e2/6), Uptime 00:45:55, Rate 115 (SM)
upstream_neighbor 52.1.1.1
Flags (0xf006c4e1) SM SPT LRCV HW FAST TAG MSDPADV RPFs
fast ports: ethe 1/16 ethe 2/9
AgeSltMsk: 00000002, FID: 0x8225, MVID: 257 , RegPkt: 44, AvgRate: 114, profile: none
Forwarding_oif: 2, Immediate_oif: 1, Blocked_oif: 1
L3 (HW) 2:
e1/16(VL57), 00:43:17/0, Flags: MJ
e2/9(VL59), 00:44:45/168, Flags: IM MJ
Blocked OIF 1:
e1/2(VL53), 00:44:00/209, Flags: IH BR
```

Multicast ECMP support

If there are multiple Equal Cost Paths between PIM routers to reach the source or the RP, multicast RPF algorithm should distribute the load across available paths to take advantage of those paths.


Figure 8 shows a topology in which R1 through R11 have IP addresses in ascending order i.e. R1 having the lowest ip address and R11 having the highest. All the routers are PIM enabled routers. The links emanating from each router are ECMP links. The existing behavior path utilization is indicated in red. With the highest IP address neighbor chosen for the ECMP paths available, the multicast cache entries get to utilize only the R1-R4-R11-SRC/RP path.

FIGURE 8 Path utilization without Multicast ECMP support



With the ECMP support turned on, the multicast entries will be distributed among the equal cost next hops as indicated in green for better utilization of the available paths.

FIGURE 9 Path utilization with Multicast ECMP support



The load distribution is achieved by distributing the multicast cache entries (*,G or S,G) to the available paths thus distributing the traffic. Two different methods are widely used to achieve this distribution.

1. Hash based - Load splitting
2. Least used path based - Load balancing

Brocade devices support the Hash based method of load distribution for multicast ECMP.

Hash based load Distribution

- Depends on a hash function to distribute the multicast cache entries.
- hash function based S, G, next-hop addresses.
- Splits the cache entries by choosing a different RPF neighbor and in turn splits the traffic.
- Load balancing is based on the distribution of the keys S, G, next-hop.
- Least disruptive as the hashing redistributes only those cache entries that are affected during link flaps.
- Some paths may not be utilized for the distribution of the multicast entries. For example, for the ECMP paths from R3 to R6, R7 and R8, only paths R3 to R7 and R3 to R8 are being utilized.

Deleting path

When an ECMP path goes down, all the multicast entries using that path get redistributed among the other available paths.

Adding path

No redistribution (default behavior without rebalance option) of the cache entries when a new path is added to the ECMP set.

- Here Optimal utilization of the paths is traded off in favor of not disturbing the existing flow.
- It also requires a full branch setup towards the source or RP of the multicast distribution tree sometimes.
- When a path flaps i.e. goes down and comes back up, the multicast entries which had been using this path would not be using this path anymore and it becomes worse if a subset of paths go down and come back up one by one, resulting in only the path(s) that didn't flap to carry all entries.

Dynamic rebalancing

- Option to rebalance the traffic immediately on a new next-hop or path addition.
- Both CONFIG and EXEC level option.
- Helps in both new next-hop and path addition and path flap cases.
- The existing flows will be disturbed with least disruption by using the hashing method.

Limitations and prerequisites

The following limitations and prerequisites apply to the configuration of ECMP path load balancing.

- The hash method is a load splitting method and hence traffic load balancing is not supported.
- S based and S,G based hashing is not supported.
- The hash method is a load splitting method and not a load balancing method and hence the load balancing effect due to load splitting the multicast entries is only a best effort and the splitting is actually based on the number of S, G flows and the number of next-hops and the actual distribution of the S,G and the next-hop addresses.
- If the rebalancing is not configured, then link flap results in sub-optimal utilization of the ECMP links.
- The number of paths supported by multicast ECMP would be the same as unicast ECMP which is 32.

Enabling Multicast ECMP

The **ip multicast-routing load-sharing** command configures the hash based distribution among the ECMP paths.

Once configured, this redistributes the existing flows among the available ECMP paths. However when a new next-hop is added to the ECMP set, the traffic will not be redistributed. Once the EXEC level command is executed, it redistributes all the IPv4 multicast flows. To configure Multicast ECMP, use this command in the configuration mode.

```
device(config-if-e10000-1/1)# ip multicast-routing load-sharing
device(config-if-e10000-1/1)# ipv6 multicast-routing load-sharing
```

To disable load distribution among ECMP paths use the no form of the command. This will revert to the default ECMP behavior of choosing the highest IP address neighbor from the ECMP set.

```
device(config-if-e10000-1/1)# no ip multicast-routing load-sharing
device(config-if-e10000-1/1)# no ipv6 multicast-routing load-sharing
```

Syntax: [no] [ip | ipv6] multicast-routing load-sharing

Enabling Rebalance

The **rebalance** option enables redistributing the load when a new next-hop is added. The redistribution is based on the hash function.

Once configured, this redistributes the existing flows among the all available ECMP paths. In addition to that whenever a new next-hop is added, some of the existing flows will be redistributed to the new path added using the newly added ECMP path. Once the EXEC level command is executed, it redistributes all the IPv4 multicast flows. To configure rebalancing, use the `rebalance` in the `multipath` command.

```
device(config-if-e10000-1/1)# ip multicast-routing load-sharing rebalance
device(config-if-e10000-1/1)# ipv6 multicast-routing load-sharing rebalance
```

To disable rebalancing when a new next-hop is added, use `no` form of the `rebalance` command.

```
device(config-if-e10000-1/1)# no ip multicast-routing load-sharing rebalance
device(config-if-e10000-1/1)# no ipv6 multicast-routing load-sharing rebalance
```

Syntax: `[no] [ip | ipv6] multicast-routing load-sharing rebalance`

Displaying ECMP paths in RPF neighbor

The `show` command displays the RPF neighbor to indicate, if the neighbor is one of the ECMP paths. If the neighbor is one of the next-hops in the PIM ECMP set then the ECMP is also included in the output at the end.

```
device(config)# show ip pim vrf eng rpf 130.50.11.10 226.10.10.1
Upstream nbr 55.55.55.55 on e4/1 - mroute entry - ecmp
```

If the group address is not specified and there are multiple paths to the unicast address mentioned then the multiple ECMP upstream neighbors will be shown.

```
device(config)# show ip pim vrf eng rpf 130.50.11.10
Upstream Nbr Phy Port
55.55.55.55 e4/1
66.55.55.55 e4/2
device(config)# show ip pim mcache load-sharing src-ip-address
device(config)# show ip pim mcache load-sharing 130.50.11.10
Source/RP Address Upstream Nbr Interface Count
130.50.11.10 55.55.55.55 e4/1 4
130.50.11.10 66.55.55.55 e4/2 3
```

If the source address is not specified, then the show output will be as follows.

```
device(config)# show ip pim mcache load-sharing
Source/RP Address Upstream Nbr Interface Count
11.11.11.11 * 192.160.1.1 e2/16 8
11.11.11.11 * 192.168.1.1 e2/15 12
130.50.11.10 55.55.55.55 e4/1 2
130.50.11.10 66.55.55.55 e4/2 7
* Indicates RP Address
```

Enabling multicast fast convergence

When you enable optimization for multicast fast convergence, each PIM join is sent immediately, which ensures faster convergence.

However, enabling fast convergence also increases the number of PIM messages on the system. In systems with very high number of mcache entries, batching of PIM messages is recommended to reduce the number of periodic messages and for faster convergence.

The following considerations apply to IP multicast fast convergence:

- PIM Sparse Mode must be enabled on MLX devices.
 - After moving to the PIM sparse mode, you may notice a black channel for a period of up to 10 seconds while joining multicast groups not in use.
1. Enter global configuration mode by issuing the **configure terminal** command.
 2. Enter the **ip multicast-routing fast-convergence** command to enable optimization for fast convergence.

Configuring Layer 3 Multicast filter for the hardware

This feature introduces the ability to program the multicast filter entries into the hardware. It is an extension of the Multicast boundary feature to filter control packets in software and data packets in hardware.

This prevents CPU punting and protects the CPU from excessive usage. The filter drop entries are programmed at the ingress TCAM. These are effective for ACL with permit clause only. You can bind only one ACL to any interface and globally under router PIM.

Limitations and pre-requisites

- It is not allowed to configure multicast filter globally and at the interface level simultaneously.
- It cannot have both software multicast boundary and multicast filter configured on a PIM enabled interface.
- Global multicast filter is not supported on non-default VRFs.

In scaled mcache or PIM interfaces scenarios, when a global multicast filter is applied, MP CPU usage may be high for approximately 1 second.

When 128 PIM interfaces are configured, the recommended upper limit for the mcache entries can be approximately 1000 (including *G and SG). When 64 PIM interfaces are configured, the number of mcache entries can be approximately 2000.

Configuring the Layer 3 Multicast filter

This command can be applied at global level under router PIM as well as on PIM enabled interface for both IPv4 and IPv6.

Configuring Layer 3 multicast filter at global level

Global Layer 3 multicast filter installs filter drop entries into hardware with Port and VLAN masked.

To configure IPv4 Layer 3 multicast filter at Global level, enter a command such as the following,

```
device(config)# router pim
device(config-pim-router)# multicast-filter MyBrocadeAccessList
```

To configure Layer 3 multicast filter for IPv6 at global level as following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# multicast-filter MyBrocadeAccessList
```

Use the **acl-spec** parameter to define the number or name identifying an access list that controls the range of group addresses affected by the Layer 3 multicast filter.

Configuring Layer 3 multicast filter at interface level

To define L3 multicast filter for PIM enabled interfaces, enter a command such as the following.

NOTE

This command is available only on CES/CER devices at the interface configuration level.

```
device(config)# interface ve 40
device(config-vif-40)# ip pim multicast-filter MyBrocadeAccessList
```

Layer 3 multicast filter can be configured for IPv6 as shown below.

```
device(config)# interface ethernet 1/2
device(config-if-e1000-1/2)# ipv6 pim multicast-filter MyBrocadeAccessList
```

The ACL, **MyBrocadeAccessList** can be configured using standard or extended ACL syntax. Some examples of how ACLs can be used to filter multicast traffic are as below.

Standard ACL to deny multicast traffic

To filter multicast traffic for group 225.1.1.2, enter the following command.

```
device(config)# access-list 10 permit host 225.1.1.2
```

Extended ACL to deny multicast traffic

To filter multicast traffic for group 226.1.1.1 from any source, enter the following command.

```
device(config)# access-list 101 permit ip any host 226.1.1.1
```

NOTE

Layer 3 multicast filter applies for ACL with permit clause only. ACL with deny clause and "permit any any" are ignored.

Layer 3 multicast filter can be applied on following PIM enabled interface types.

- Regular port: Multicast filter drop entries will be programmed with p,v set.
- VE: Multicast drop entry will be programmed with the VLAN set and port set to wild card. For VE type, CAM entry will be installed on all LP's and on all the PPCR's irrespective of port present on a PPCR.
- LAG: Multicast drop entry will be programmed with p,v set for all the configured LAG members. LAG members up or down event will not be handled, only the members addition and deletion will be handled and programmed accordingly.

- GRE: Multicast drop entry will be programmed with p,v masked as long as the incoming interface is identified as a tunnel interface

Displaying Multicast filter for the hardware

This command displays the multicast filters attached to egress ports.

This command can be applied under router PIM as well as PIM enabled interface at MP and LP.

Proper validation checks prevent configuring another Multicast filter ACL on PIM enabled interface or globally without un-configuring the previous Multicast filter.

1. Return to global configuration mode.
2. Layer 3 multicast filter output at MP, when Global filter is applied.

```
device(config)# show ip pim multicast-filter
-----+-----+-----+
Interface      | LAG      | Multicast Filter
                | Member   |
-----+-----+-----+
*              | -        | (1.1.1.1, 239.1.1.1)
*              | -        | (2.2.2.2, 225.1.1.0/24)
```

Show output at LP for global filter.

```
device(config)# show ip pim multicast-filter
-----+-----+-----+-----+-----+-----+-----+
Interface|LAG Member |port |vlan | Multicast Filter |CAM Index| ProgTM
-----+-----+-----+-----+-----+-----+-----+
*        | -         | *   | *   | 1.1.1.1, 239.1.1.1 | 0x343   | 22:01:33
*        | -         | *   | *   | *, 234.1.1.1      | *,      | 0x344   22:01:33
```

Layer 3 multicast filter output at MP, when Interface filter is applied.

```
device(config)# show ip pim multicast-filter ve 100
-----+-----+-----+
Interface      | LAG      | Multicast Filter
                | Member   |
-----+-----+-----+
ve100          | -        | (1.1.1.1, 239.1.1.1)
```

Layer 3 multicast filter output at LP when Interface filter is applied.

```
device(config)# show ip pim multicast-filter
-----+-----+-----+-----+-----+-----+-----+
Interface      |LAG Member |port |vlan | Multicast Filter          |CAM Index|ProgTM
-----+-----+-----+-----+-----+-----+-----+
ve100          | -         | *   | 100 | 1.1.1.1, 239.1.1.1      | 0x343   | 22:01:33
ve102          | -         | *   | 100 | *, 234.1.1.1           | *,      | 0x344   22:01:33
e1/13         | -         | 142 | 1   | *, 226.1.1.1           | 0x355   | 22:01:33
Tr1(e1/1)     | e1/1     | 155 | 1   | *, 225.1.1.1           | *,      | 0x356   22:01:33
Tr1           | e1/4     | 156 | 1   | 2.2.2.2, 225.1.1.0/240x357 | 22:01:33
Tn1           | -         | *   | *   | *, 227.1.1.2           | 0x358   | 22:01:33
```

The output displays the multicast filters globally and for individual interfaces respectively.

PIM Dense

NOTE

This section describes the "dense" mode of PIM, described in RFC 3973. Refer to [PIM Sparse](#) on page 80 for information about PIM Sparse.

PIM was introduced to simplify some of the complexity of the routing protocol at the cost of additional overhead tied with a greater replication of forwarded multicast packets. PIM builds source-routed multicast delivery trees and employs reverse path check when forwarding multicast packets.

There are two modes in which PIM operates: Dense and Sparse. The Dense Mode is suitable for densely populated multicast groups, primarily in the LAN environment. The Sparse Mode is suitable for sparsely populated multicast groups with the focus on WAN.

PIM primarily uses the IP routing table instead of maintaining its own, thereby being routing protocol independent.

Initiating PIM multicasts on a network

Once PIM is enabled on each device, a network user can begin a video conference multicast from the server on R1 as shown in [Pruning a multicast tree](#) on page 65. When a multicast packet is received on a PIM-capable device interface, the interface checks its IP routing table to determine whether the interface that received the message provides the shortest path back to the source. If the interface does provide the shortest path back to the source, the multicast packet is then forwarded to all neighboring PIM devices. Otherwise, the multicast packet is discarded and a prune message is sent back upstream.

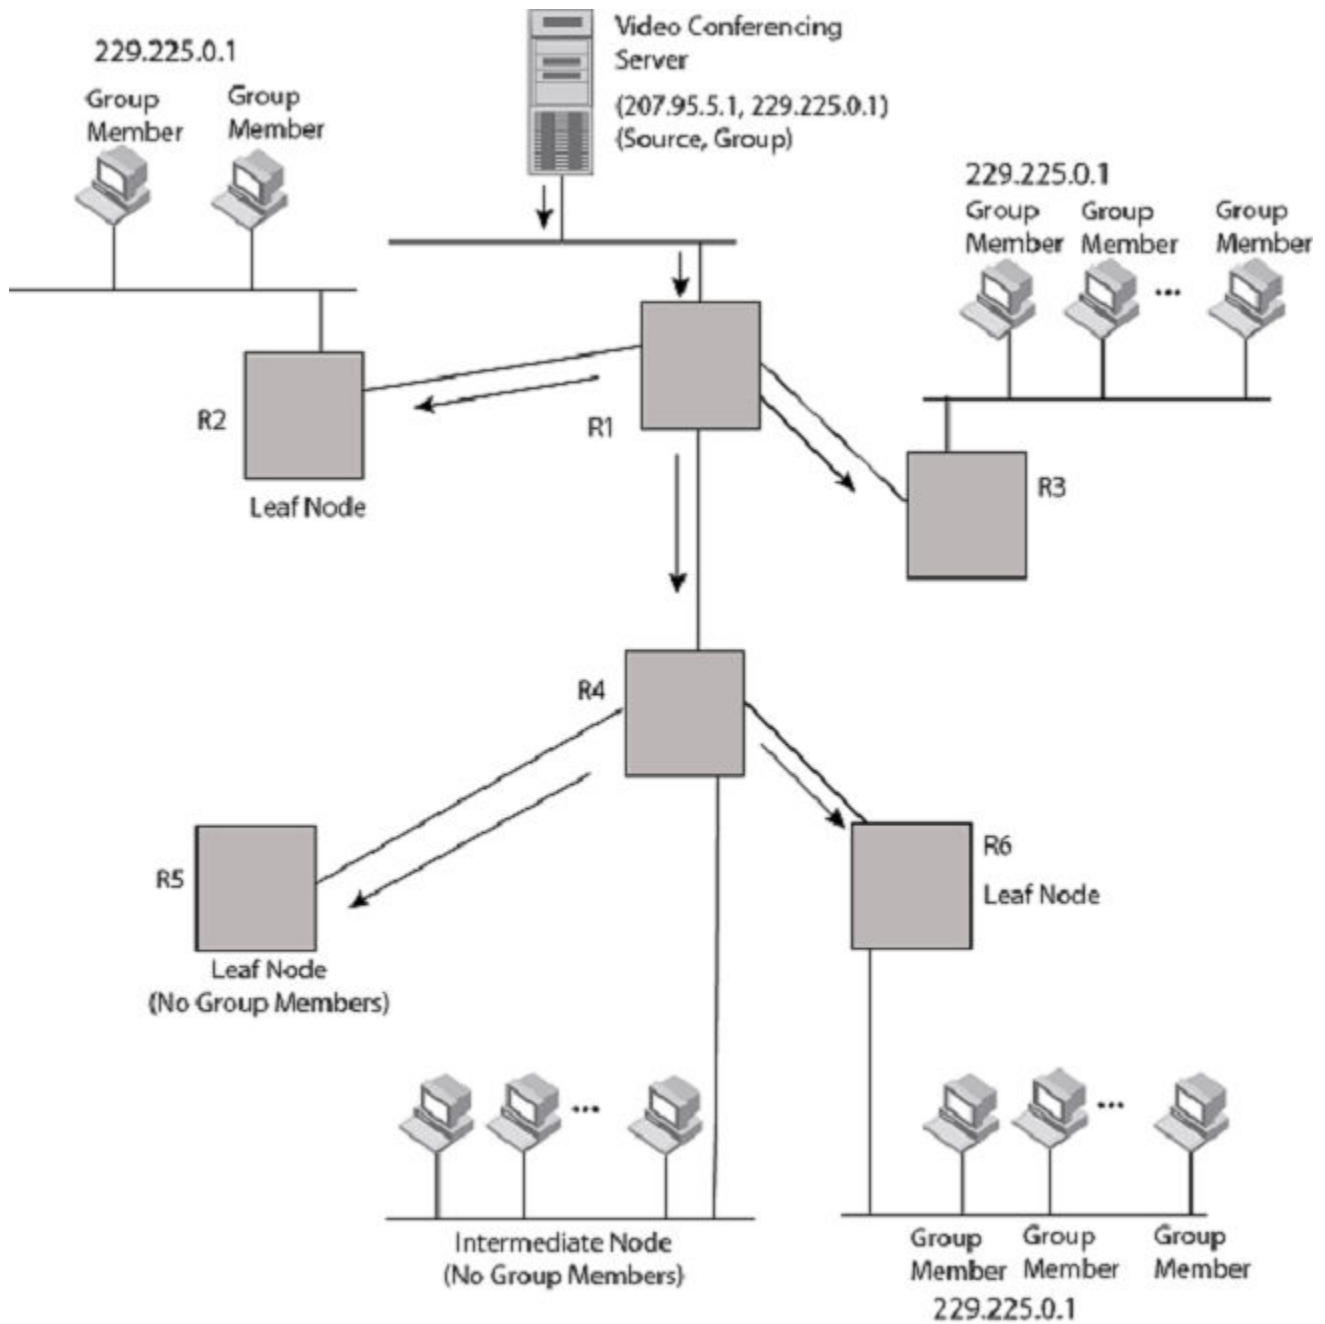
In [Pruning a multicast tree](#) on page 65, the root node (R1) is forwarding multicast packets for group 229.225.0.1, which it receives from the server, to its downstream nodes, R2, R3, and R4. Device R4 is an intermediate device with R5 and R6 as its downstream devices. Because R5 and R6 have no downstream interfaces, they are leaf nodes. The receivers in this example are those workstations that are resident on devices R2, R3, and R6.

Pruning a multicast tree

As multicast packets reach these leaf devices, the devices check their IGMP databases for the group. If the group is not in the IGMP database of the device, the device discards the packet and sends a prune message to the upstream device. The device that discarded the packet also maintains the prune state for the source, group (S,G) pair. The branch is then pruned (removed) from the multicast tree. No further multicast packets for that specific (S,G) pair will be received from that upstream device until the prune state expires. You can configure the PIM Prune Timer (the length of time that a prune state is considered valid).

For example, in the following figure, the sender with address 207.95.5.1 is sending multicast packets to the group 229.225.0.1. If a PIM device receives any groups other than that group, the device discards the group and sends a prune message to the upstream PIM device.


FIGURE 10 Transmission of multicast packets from the source to host group members



In the following figure, Device R5 is a leaf node with no group members in its IGMP database. Therefore, the device must be pruned from the multicast tree. R5 sends a prune message upstream to its neighbor device R4 to remove itself from the multicast delivery tree and install a prune state. Device R5 will not receive any further multicast traffic until the prune age interval expires.

When a node on the multicast delivery tree has all of its downstream branches (downstream interfaces) in the prune state, a prune message is sent upstream. In the case of R4, if both R5 and R6 are in a prune state at the same time, R4 becomes a leaf node with no downstream interfaces and sends a prune message to R1. With R4 in a prune state, the resulting multicast delivery tree would consist only of leaf nodes R2 and R3.

FIGURE 11 Pruning leaf nodes from a multicast tree



Grafts to a multicast tree

A PIM device restores pruned branches to a multicast tree by sending graft messages towards the upstream device. Graft messages start at the leaf node and travel up the tree, first sending the message to its neighbor upstream device.

In the example above, if a new 229.255.0.1 group member joins on device R6, which was previously pruned, a graft is sent upstream to R4. Since the forwarding state for this entry is in a prune state, R4 sends a graft to R1. Once R4 has joined the tree, R4 along with R6 once again receive multicast packets.

Prune and graft messages are continuously used to maintain the multicast delivery tree. No configuration is required on your part.

PIM DM versions

The Extreme device supports PIM DM V1 and V2. The default is V2. You can specify the version on an individual interface basis.

The primary difference between PIM DM V1 and V2 is the methods the protocols use for messaging:

- PIM DM V1 - uses the IGMP to send messages.
- PIM DM V2 - sends messages to the multicast address 224.0.0.13 (ALL-PIM-ROUTERS) with protocol number 103.

The CLI commands for configuring and managing PIM DM are the same for V1 and V2. The only difference is the command you use to enable the protocol on an interface.

NOTE

If you want to continue to use PIM DM V1 on an interface, you must change the version, then save the configuration.

NOTE

The note above does not mean you can run different PIM versions on devices that are connected to each other. The devices must run the same version of PIM. If you want to connect a Extreme device running PIM to a device that is running PIM V1, you must change the PIM version on the Extreme device to V1 (or change the version on the device to V2, if supported).

Configuring PIM DM

NOTE

This section describes how to configure the "dense" mode of PIM, described in RFC 1075. Refer to [Configuring PIM Sparse](#) on page 82 for information about configuring PIM Sparse.

Enabling PIM on the device and an interface

By default, PIM is disabled. To enable PIM:

- Enable the feature globally.
- Configure the IP interfaces that will use PIM.
- Enable PIM locally on the ports that have the IP interfaces you configured for PIM.

Suppose you want to initiate the use of desktop video for fellow users on a sprawling campus network. All destination workstations have the appropriate hardware and software but the devices that connect the various buildings need to be configured to support PIM multicasts from the designated video conference server as shown in [Pruning a multicast tree](#) on page 65.

PIM is enabled on each of the devices shown in [Pruning a multicast tree](#) on page 65, on which multicasts are expected. You can enable PIM on each device independently or remotely from one of the devices with a Telnet connection. Follow the same steps for each device. All changes are dynamic.

Globally enabling and disabling PIM

To globally enable PIM, enter the following command.

```
device(config)# router pim
```

Syntax: [no] router pim

NOTE

When PIM routing is enabled, the line rate for receive traffic is reduced by about 5%. The reduction occurs due to overhead from the VLAN multicasting feature, which PIM routing uses. This behavior is normal and does not indicate a problem with the device.

The **[no] router pim** command behaves in the following manner:

- Entering **router pim** command to enable PIM does not require a software reload.
- Entering a **no router pim** command removes all configuration for PIM multicast on a device (**router pim** level) only.

Enabling PIM for a specified VRF

To enable PIM for the VRF named "blue", use the following commands.

```
device(config)# router pim vrf blue
```

Syntax: **[no] router pim [vrf vrf-name]**

The **vrf** parameter allows you to configure PIM (PIM-DM and PIM-SM) on the virtual routing instance (VRF) specified by the **vrf-name** variable. All PIM parameters available for the default device instance are configurable for a VRF-based PIM instance.

The **[no] router pim vrf** command behaves in the following manner:

- Entering the **router pim vrf** command to enable PIM does not require a software reload.
- Entering a **no router pim vrf** command removes all configuration for PIM multicast on the specified VRF.

Enabling a PIM version

To enable PIM on an interface, globally enable PIM, then enable PIM on interface 3, enter the following commands.

```
device(config)# router pim
device(config)# int e 1/3
device(config-if-e10000-1/3)# ip address 207.95.5.1/24
device(config-if-e10000-1/3)# ip pim
device(config-if-e10000-1/3)# write memory
device(config-if-e10000-1/3)# end
```

Syntax: **[no] ip pim [version 1 | 2]**

The **version 1 | 2** parameter specifies the PIM DM version. The default version is 2.

If you have enabled PIM version 1 but need to enable version 2 instead, enter either of the following commands at the configuration level for the interface.

```
device(config-if-e10000-1/1)# ip pim version 2
device(config-if-e10000-1/1)# no ip pim version 1
```

To disable PIM DM on the interface, enter the following command.

```
device(config-if-e10000-1/1)# no ip pim
```

Modifying PIM global parameters

PIM global parameters come with preset values. The defaults work well in most networks, but you can modify the following parameters if necessary:

- Neighbor timeout
- Hello timer
- Prune timer

- Prune wait timer
- Graft retransmit timer
- Inactivity timer

Modifying neighbor timeout

Neighbor timeout is the interval after which a PIM device will consider a neighbor to be absent. Absence of PIM hello messages from a neighboring device indicates that a neighbor is not present.

The interval can be set between 3 and 65535 seconds, and it should not be less than 3.5 times the hello timer value. The default value is 105 seconds.

To apply a PIM neighbor timeout value of 360 seconds to all ports on the device operating with PIM, enter the following.

```
device(config)# router pim
device(config-pim-router)# nbr-timeout 360
```

Syntax: **[no]** **nbr-timeout** *seconds*

The default is 105 seconds. The range is 3-65535 seconds.

Modifying hello timer

This parameter defines the interval at which periodic hellos are sent out PIM interfaces. Devices use hello messages to inform neighboring devices of their presence. The interval can be set between 1 and 3600 seconds, and the default rate is 30 seconds.

To apply a PIM hello timer of 120 seconds to all ports on the device operating with PIM, enter the following.

```
device(config)# router pim
device(config-pim-router)# hello-timer 120
```

Syntax: **[no]** **hello-timer** *1-3600*

The default is 30 seconds.

Modifying prune timer

This parameter defines how long a PIM device will maintain a prune state for a forwarding entry.

The first received multicast interface is forwarded to all other PIM interfaces on the device. If there is no presence of groups on that interface, the leaf node sends a prune message upstream and stores a prune state. This prune state travels up the tree and installs a prune state.

A prune state is maintained until the prune timer expires or a graft message is received for the forwarding entry. The default value is 180 seconds.

To set the PIM prune timer to 90, enter the following.

```
device(config)# router pim
device(config-pim-router)# prune-timer 90
```

Syntax: **[no]** **prune-timer** *seconds*

The default is 180 seconds. The range is 60-3600 seconds.

NOTE

The prune timer should always be configured in multiples of 30 within the range of 60–3600. When you assign any other value, the prune timer will be changed to the next upper limit in the running configuration. For example, if you assign a prune timer for 70, which is between 60 and 90, the prune timer will be set to 90 in the running configuration.

Modifying the prune wait timer

The **prune-wait** command allows you to configure the amount of time a PIM device will wait before stopping traffic to neighbor devices that do not want the traffic. The value can be from zero to fifteen seconds. The default is three seconds. A smaller prune wait value reduces flooding of unwanted traffic.

A prune wait value of zero causes the PIM device to stop traffic immediately upon receiving a prune message. If there are two or more neighbors on the physical port, then the prune-wait command should not be used because one neighbor may send a prune message while the other sends a join message at the same time, or within less than three seconds.

To set the prune wait time to zero, enter the following commands.

```
device(config)#router pim
device(config-pim-router)#prune-wait 0
```

Syntax: [no] prune-wait seconds

The **seconds** can be 0 - 30. A value of 0 causes the PIM device to stop traffic immediately upon receiving a prune message. The default is 3 seconds.

To view the currently configured prune wait time, enter the **show ip pim dense** command as described in [Displaying basic PIM Dense configuration information](#) on page 73.

Modifying graft retransmit timer

The graft retransmit timer defines the interval between the transmission of graft messages.

A graft message is sent by a device to cancel a prune state. When a device receives a graft message, the device responds with a Graft Ack (acknowledge) message. If this Graft Ack message is lost, the device that sent the graft message will resend it.

To change the graft retransmit timer from the default of 180 to 90 seconds, enter the following.

```
device(config)# router pim
device(config-pim-router)# graft-retransmit-timer 90
```

Syntax: [no] graft-retransmit-timer seconds

The default is 180 seconds. The range is from 60–3600 seconds.

Modifying inactivity timer

The device deletes a forwarding entry if the entry is not used to send multicast packets. The PIM inactivity timer defines how long a forwarding entry can remain unused before the device deletes it.

To apply a PIM inactivity timer of 90 seconds to all PIM interfaces, enter the following.

```
device(config)# router pim
device(config-pim-router)# inactivity-timer 90
```

Syntax: [no] inactivity-timer seconds

The default is 180 seconds. The range is from 10–3600 seconds.

Selection of shortest path back to source

By default, when a multicast packet is received on a PIM-capable interface in a multi-path topology, the interface checks its IP routing table to determine the shortest path back to the source. If the alternate paths have the same cost, the first alternate path in the table is picked as the path back to the source. For example, in the following example, the first four routes have the same cost back to the source. However, 137.80.127.3 is chosen as the path to the source since it is the first one on the list. The device rejects traffic from any port other than Port V11 on which 137.80.127.3 resides

```
Total number of IP routes: 19
Type Codes - B:BGP D:Connected I:ISIS S:Static R:RIP O:OSPF Cost - Dist/Metric
Destination      Gateway          Port      Cost      Type
..
172.17.41.4      137.80.127.3    v11       2          O
172.17.41.4      137.80.126.3    v10       2          O
172.17.41.4      137.80.129.1    v13       2          O
172.17.41.4      137.80.128.3    v12       2          O
172.17.41.8      0.0.0.0         1/2       1          D
```

Failover time in a multi-path topology

When a port in a multi-path topology fails, multicast devices, depending on the routing protocol being used, take a few seconds to establish a new path, if the failed port is the input port of the downstream device.

Modifying the TTL threshold

The TTL threshold defines the minimum value required in a packet for it to be forwarded OUT of the interface AFTER the TTL has been decremented.

For example, if the TTL for an interface is set at 10, only those packets that enter with a TTL value of 11 or more are forwarded through the TTL-10 interface. With a default TTL threshold of 1, only packets ingressing with a TTL of 2 or greater are forwarded. The TTL threshold only applies to routed interfaces and is ignored by switched interfaces. Possible TTL values are 1 to 64. The default TTL value is 1.

To configure a TTL of 45, enter a command such as the following.

```
device(config-if-e10000-3/24)# ip pim ttl-threshold 45
```

Syntax: [no] ip pim ttl-threshold 1-64

Configuring a DR priority

The DR priority option lets you give preference to a particular device in the DR election process by assigning it a numerically higher DR priority. This value can be set for IPv4 and IPv6 interfaces. To set a DR priority higher than the default value of 1, use the **ip pim dr-priority** command as shown:

For IPv4.

```
device(config-if-e10000-3/24)# ip pim dr-priority 50
```

For IPv6.

```
device(config-if-e10000-3/24)# ipv6 pim dr-priority 50
```

Syntax: [no] ip pim dr-priority *priority-value*

Syntax: [no] ipv6 pim dr-priority *priority-value*

The **priority-value** variable is the value that you want to set for the DR priority. Optional values are: 0 - 65535. The default value is 1.

The **no** option removes the command and sets the DR priority back to the default value of 1.

The following information may be useful for troubleshooting.

1. If more than one device has the same DR priority on a subnet (as in the case of default DR priority on all), the device with the numerically highest IP address on that subnet is elected as the DR.
2. The DR priority information is used in the DR election ONLY IF ALL the PIM devices connected to the subnet support the DR priority option. If there is at least one PIM device on the subnet that does not support this option, then the DR election falls back to the backwards compatibility mode in which the device with the numerically highest IP address on the subnet is declared the DR regardless of the DR priority values.

Displaying basic PIM Dense configuration information

To display PIM Dense configuration information, enter the following command at any CLI level.

```
device(config)# show ip pim dense
Global PIM Dense Mode Settings
  Maximum Mcache           : 0           Current Count           : 500
  Hello interval          : 30           Neighbor timeout        : 105
  Join/Prune interval     : 60           Inactivity interval     : 180
  Hardware Drop Enabled   : Yes         Prune Wait Interval    : 3
  Graft Retransmit interval : 180       Prune Age              : 180
  Route Precedence        : mc-non-default mc-default uc-non-default uc-default
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Interface|Local      |Ver |St | Designated Router |TTL|Multicast| VRF | DR |Override|
         |Address   |   |   | Address          |Port|Thr|Boundary | Prio|Interval|
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
e2/2     |103.103.1.1|DMv2|Ena| 103.103.1.2    |2/2| 1|None| default| 1 |3000ms
v102    |102.1.1.2  |DMv2|Ena| Itself          |   | 1|None| default| 1 |3000ms
v107    |107.1.1.1  |DMv2|Ena| Itself          |   | 1|None| default| 1 |3000ms
v109    |109.1.1.1  |DMv2|Dis| Itself          |   | 1|None| default| 1 |3000ms
Total Number of Interfaces : 4
device(config)#
```

Syntax: `show ip pim [vrf vrf-name] dense`

The **vrf** option allows you to display PIM dense configuration information for the VRF instance identified by the *vrf-name* variable.

This display shows the following information.

This field...	Displays...
Maximum Mcache	The maximum number multicast cache entries allowed on the device.
Current Count	The number of multicast cache entries currently used.
Hello interval	How frequently the device sends hello messages out the PIM dense interfaces.
Neighbor timeout	The interval after which a PIM device will consider a neighbor to be absent.
Graft or Retransmit interval	How interval between the transmission of graft messages.
Inactivity interval	How long a forwarding entry can remain unused before the device deletes it.
Join or Prune interval	How long a PIM device will maintain a prune state for a forwarding entry.
Prune Age	The number of packets the device sends using the path through the RP before switching to using the SPT path.
Hardware Drop Enabled	Displays Yes if the Passive Multicast Route Insertion feature is enable and No if it is not.
Prune Wait Interval	The amount of time a PIM device waits before stopping traffic to neighbor devices that do not want the traffic. The value can be from zero to three seconds. The default is three seconds.

This field...	Displays...
Route Precedence	The route precedence configured to control the selection of routes based on the route types. There are four different types of routes: <ul style="list-style-type: none"> • Non-default route from the mRTM • Default route from the mRTM • Non-default route from the uRTM • Default route from the uRTM
Interface	The type of interface and the interface number.
Local Address	Indicates the IP address configured on the port or virtual interface.
Mode	Either DM for Dense Mode or SM for Sparse Mode.
Ver	The version of PIM Dense mode.
Designated Router	The IP address and port for the PIM designated device.
TTL Threshold	The minimum value required in a packet for it to be forwarded out of the interface.
Multicast Boundary	The multicast boundary enabled (if any) for PIM interfaces.
VRF	If the PIM Dense instance is configured within a VRF, this field will contain the name.
DR Prio	The priority of the designated device.

Displaying all multicast cache entries in a pruned state

Use the following command to display all multicast cache entries that are currently in a pruned state and have not yet aged out.

```
device(config)# show ip pim prune
 1 (104.1.1.2 231.0.1.1):
 e2/2,2/2(150)
 2 (108.1.1.100 231.0.1.1):
 e2/2,2/2(150)
 3 (104.1.1.2 231.0.1.2):
 e2/2,2/2(150)
 4 (108.1.1.100 231.0.1.2):
 e2/2,2/2(150)
 5 (108.1.1.100 231.0.1.3):
 e2/2,2/2(150)
 6 (104.1.1.2 231.0.1.4):
 e2/2,2/2(150)
 7 (108.1.1.100 231.0.1.4):
 e2/2,2/2(150)
 8 (104.1.1.2 231.0.1.5):
 e2/2,2/2(150)
 9 (108.1.1.100 231.0.1.5):
 e2/2,2/2(150)
Total Prune entries: 9
```

Syntax: `show ip pim [vrf vrf-name] prune`

Displaying all multicast cache entries

You can use the following command to display all multicast cache entries.

```
device(config)# show ip pim mcache
Total entries in mcache: 234
 1 (104.1.1.2, 231.0.1.1) in v102 (tag e1/19), Uptime 00:02:15 Rate 0
 upstream neighbor=102.1.1.1
 fast ports
 Prunes: e2/2,2/2(150)
```

```
Flags (0x300004c9)
  sm=0 ssm=0 hw=1 fast=1 slow=0 leaf=0 prun=1 tag=0 needRte=0 msdp_adv=0
AgeSltMsk=00000001, FID: 0x8000 MVID: NotReq, AvgRate 0 profile: none
```

Syntax: `show ip pim mcache [source-address | group-address | counts | dense | fid fid-id | g_entries | mvid mvid | receiver | sg_entries | sparse | ssm]`

The *source-address* parameter selects the multicast cache source address.

The *group-address* parameter selects the multicast cache group address.

The **counts** keyword indicates the count of entries.

The **dense** keyword displays only the PIM Dense Mode entries.

The *fid-id* variable allows you to display all entries that match a specified fid .

The **g_entries** keyword displays only the (*, G) entries .

The *mvid* variable allows you to display all entries that match a specified mvid.

The **receiver** keyword allows you to display all entries that egress a specified interface.

The **sg_entries** keyword displays only the (S, G) entries .

The **sparse** keyword displays only the PIM Sparse Mode entries.

The **ssm** keyword displays only the SSM entries.

NOTE

In NetIron CES and NetIron CER devices, the hardware decides the forwarding port for the LAG. The software does not have the capability to guess which port the packet will go on. So software fwd port in the OIF may be different than the actual fwd port in hardware on a trunk link.

TABLE 8 Output fields from the `show ip pim mcache` command

Field	Description
Total entries in mcache	Shows the total number of PIM mcache entries
MJ	Membership Join
MI	Membership Include
ME	Membership Exclude - Legend for the mcache entry printed once per page, it gives the explanation of each of the flags used in the entry.
BR	Blocked RPT
BA	Blocked Assert
BF	Blocked Filter
BI	Blocked IIF
Uptime	Shows the software entry uptime. This field is displayed on both the MP and the LP modules.
Rate	Shows the total number of packets per second that have been forwarded using the hardware programmed forwarding entry (the (S,G) entry programmed in hardware or (*,G) entries if (*,G) based forwarding is enabled). The rate is displayed for all entries when the fwd_fast flag is set on the MP module.
upstream neighbor	Shows the upstream neighbor for the Source/RP based on the type of entry. For (*,G) it shows the upstream neighbor towards the RP. For (S,G) entries it shows the upstream neighbor towards the source.

TABLE 8 Output fields from the `show ip pim mcache` command (continued)

Field	Description
Flags	<p>Flags Represent Entry flags in hex format in the braces. And indicates the meaning of the flags set in abbreviated string whose explanations are as below. Only shows the flags which are set.</p> <p>SM - Shows If the entry is created by PIM Sparse Mode</p> <p>DM - Shows If DM mode entry is enabled</p> <p>SSM - Shows If the SSM mode entry is enabled</p> <p>RPT - Shows If the entry is on the Rendezvous Point (RP)</p> <p>SPT - Shows If the entry is on the source tree</p> <p>LSRC - Shows If the source is in a directly-connected interface</p> <p>LRcv - Shows If the receiver is directly connected to the router</p> <p>REG - if the data registration is in progress</p> <p>L2REG - if the source is directly connected to the router</p> <p>REGSUPP - if the register suppression timer is running</p> <p>RegProbe</p> <p>HW - Shows If the candidate for hardware forwarding is enabled</p> <p>FAST - Shows If the resources are allocated for hardware forwarding</p> <p>TAG - Shows If there is a need for allocating entries from the replication table</p> <p>MSDPADV - Shows If RP is responsible for the source and must be advertised to its peers.</p> <p>NEEDRTE - Shows If there is no route to the source and RP is available</p> <p>PRUNE - Shows If PIM DM Prune to upstream is required</p>
RP	Show the IP address of the RP.
fast ports	Shows forwarding port mask.
AgeSlitMsk	Shows the slot number on which MP expects ingress traffic.
FID, MVID	Shows the resources that are allocated for forwarding the entry.
RegPkt	Shows the number of packets forwarded due to the Register decapsulation. This field is displayed only on the MP module. This field displays only those entries for which the device is the RP. However, for a PIM DM entry the RegPkt field is not displayed for the (S,G) entries on the MP module.
AvgRate	Shows the average rate of packets ingressing for this entry over a 30 second period. This field is displayed only on the MP module for all entries that are hardware programmed (the fwd_fast flag is set on the MP module).
ProgTm	Shows the hardware uptime of an entry. This field is only displayed on the LP module for all entries.
SwFwd	Shows the number of packets that use the software forwarding entry to forward packets. This field is displayed for all entries on the LP module.
Profile	Shows the Profile ID associated with the Stream.
Number of matching entries	Shows the total number of mcache entries matching a particular multicast filter specified.

TABLE 8 Output fields from the `show ip pim mcache` command (continued)

Field	Description
Outgoing interfaces Section	This section consists of three parts. L3 OIFs, L2OIFs and Blocked OIFs. And each section has Format of L3/L2/Blocked followed by (HW/SW) followed by count of the number of OIF in each section. Additionally, each section displays the OIFs one per line. And shows the OIF in the format eth/Tr(Vlan) followed by uptime/expiry time, followed by the Flags associated with each OIF.
L3	Show whether the traffic is routed out of the interface.
L2	Show whether the traffic is switched out of the interface.
HW	Show whether the entry is hardware forwarded.
SW	Show whether the entry is software forwarded
Eth/Tr(VL1)	Shows the outgoing interface on the specified VLAN.
Flags (explanation of flags in the OIF section)	Shows the flags set in each of the Outgoing interface in abbreviated string format whose explanations are as below. Legend of this shown at the top of each entry IM - Immediate IH - Inherited MJ - Membership Join MI - Membership Include ME - Membership Exclude BR - Blocked due to SG RPT BA - Blocked due to Assert BF - Blocked due to Filter BI - Blocked IIF (Incoming interface) matches OIF

You can use the following command to filter the output to display only entries that egress port ethernet 1/1.

```
device#show ip pim mcache receiver ethernet 1/1
```

You can use the following command to filter the output to display only the Source Specific Multicast (SSM) routes in the mcache.

```
device#show ip pim mcache ssm
```

You can use the following command to filter the output to display only the Sparse Mode routes in the mcache.

```
device#show ip pim mcache sparse
```

You can use the following command to filter the output to display only the Dense Mode routes in the mcache.

```
device#show ip pim mcache dense
```

You can use the following command to filter the output to display only the entries matching a specific source.

```
device#show ip pim mcache 1.1.1.1
```

You can use the following command to filter the output to display only the entries matching a specific group.

```
device#show ip pim mcache 239.1.1.1
```

Displaying information across VRFs

Use the following command to display information across all active VRFs.

```
device#show ip pim all-vrf ?
  bsr          Bootstrap router
  interface    PIM interface
  neighbor     PIM neighbor states
  resource     PIM resources
  rp-set       List of rendezvous point (RP) candidates
```

Multicast PIM neighbor filter

When two PIM enabled neighbor routers exchange Hello packets at regular intervals, they become PIM neighbors by default. This feature enables a router to have more control on which routers can be its neighbors by specifying an IP access-list. The access list denies PIM Hello packets from the source it wants to filter, thereby preventing that router.

Feature requirement

This feature requires the assistance of the already existing ACL feature to filter out the traffic, in the form of PIM Hello packets, from unwanted PIM neighbors.

FIGURE 12 Multicast PIM filter topology




TABLE 9 Configurations for devices running multicast PIM filters

Device 1	Device 2
interface ethernet 1/1	interface ethernet 1/2
enable	enable
ip address 10.0.0.1/24	ip address 10.0.0.2/24
ip pim-sparse	ip pim-sparse
access-list 10 deny host 10.0.0.2	
access-list 10 permit any	

Limitation and prerequisites

Following are the behavior of ACL which affect PIM-neighbor filtering process.

1. All ACLs have an implicit **deny all** clause unless overridden by an explicit **permit all** clause. Brocade devices override it for an ACL without any clauses. When you apply an empty ACL to an interface, one without any clauses, it allows all traffic on the interface to pass through without filtering.
2. There are no checks to validate if an ACL applied to an interface already exists. A non-existent ACL is considered to be equivalent to an empty ACL without any filtering capacity. A warning message is issued on the console.
3. Only one ACL can be bound to any interface.

4. It supports maximum of 128 PIM neighbor filters for both IPv4 and IPv6.

Configuring neighbor filtering

When you apply the neighbor filter on the router, use the **access-list** command to define an access-list that defines the routers you want to permit and deny to participate in PIM. The CLI options available to bind an ACL as neighbor filter are as follows:

```
<1-99>          Standard IP access list
  ASCII string   Access List Name
<100-199>      Extended IP access list
```

Defining an IPv4 standard ACL:

```
device(config)# access-list 10 deny host 10.10.10.2
device(config)# access-list 10 permit any
```

For IPv6 ACL:

```
device(config)# ipv6 access-list f10
device(config-ipv6-access-list f10)# deny ipv6 host fe80::102 any
device(config-ipv6-access-list f10)# permit ipv6 any any
```

Here fe80::102 is Link Local address of that interface.

```
device(config-if-1/3)# ip|ipv6 pim ?
```

neighbor-filter filters neighbor to participate in PIM.

Syntax: [no] [ip | ipv6] pim

Use the **pim neighbor-filter** command on an interface to filter the neighbor routers.

```
device(config)# interface ethernet 1/3
device(config-if-e1000-1/3)# ip pim neighbor-filter 10
device(config-if-e1000-1/3)# ipv6 pim neighbor-filter f10
```

This command prevents the host 10.10.10.2 as specified in access-list from becoming a PIM neighbor on interface eth 1/3.

Syntax: [no] [ip | ipv6] pim neighbor-filter *acl name*

This command applies an ACL as a rule for neighbor-filter to an interface. The ACL can either be named or numbered (standard, extended) for IPv4 and named for IPv6.

no version of the command removes the neighbor filtering applied on that interface, if any. It is not mandatory to provide Access-list name or number as at most only one ACL can be applied per interface.

Displaying show command

These commands will display the configuration.

```
device(config)#show run
!
interface ethernet 1/3
 enable
 ip address 10.10.10.1/24
 ip pim-sparse
 ip pim neighbor-filter 10
 ipv6 address a100:1111::9/64
 ipv6 pim-sparse
 ipv6 pim neighbor-filter f10
.....
```

The **show ip|ipv6 pim interface** now incorporates the Filter ACL information field as well.

On turning on **debug ip pim nbr-change** on a PIM-enabled interface, the debug messages analyze that Hello packets are received from the neighbor. On applying the neighbor-filter for the sender on that interface, the Hello packets are still received but are instantly dropped thereby preventing the sender host to become a PIM neighbor on that interface.

```
device(config)#show ip pim interface
```

```
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Interface|Local  |Ver|St|Router      |TTL|Multicast| Filter|VRF  | DR  |Override
  |Address|  |  |Address Port|Thr|Boundary | ACL  |    | Prio|Interval
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
e1/3     3.3.3.1 DMv2 Ena Itself      1 None      10  default 1  3000ms
e1/2     2.2.2.1 DMv2 Ena 2.2.2.2 1/2 1 None      None default 1  3000ms
Total Number of Interfaces: 2
```


PIM Sparse

Extreme devices support Protocol Independent Multicast (PIM) Sparse version 2. PIM Sparse provides multicasting that is especially suitable for widely distributed multicast environments. The Extreme implementation is based on RFC 2362.

In a PIM Sparse network, a PIM Sparse device that is connected to a host that wants to receive information for a multicast group must explicitly send a join request on behalf of the receiver (host).

PIM Sparse devices are organized into domains. A PIM Sparse domain is a contiguous set of devices that all implement PIM and are configured to operate within a common boundary. [Figure 13](#) shows a simple example of a PIM Sparse domain. This example shows three devices configured as PIM Sparse devices. The configuration is described in detail in the following figure.

FIGURE 13 Example PIM Sparse domain



PIM Sparse device types

Devices that are configured with PIM Sparse interfaces also can be configured to fill one or more of the following roles:

- **PMBR** - A PIM device that has some interfaces within the PIM domain and other interface outside the PIM domain. PMBRs connect the PIM domain to the Internet.
- **BSR** - The Bootstrap Router (BSR) distributes RP information to the other PIM Sparse devices within the domain. Each PIM Sparse domain has one active BSR. For redundancy, you can configure ports on multiple devices as candidate BSRs. The PIM Sparse protocol uses an election process to select one of the candidate BSRs as the BSR for the domain. The BSR with the highest BSR priority (a user-configurable parameter) is elected. If the priorities result in a tie, then the candidate BSR interface with the highest IP address is elected. In the example in [PIM Sparse](#) on page 80, PIM Sparse device B is the BSR. Port 2/2 is configured as a candidate BSR.
- **RP** - The RP is the meeting point for PIM Sparse sources and receivers. A PIM Sparse domain can have multiple RPs, but each PIM Sparse multicast group address can have only one active RP. PIM Sparse devices learn the addresses of RPs and the groups for which they are responsible from messages that the BSR sends to each of the PIM Sparse devices. In the example in [PIM Sparse](#) on page 80, PIM Sparse device B is the RP. Port 2/2 is configured as a candidate Rendezvous Point (RP). To enhance overall network performance, the Extreme device uses the RP to forward only the first packet from a group source to the group's receivers. After the first packet, the Extreme device calculates the shortest path between the receiver and source (the

Shortest Path Tree, or SPT) and uses the SPT for subsequent packets from the source to the receiver. The Extreme device calculates a separate SPT for each source-receiver pair.

NOTE

It is recommended that you configure the same ports as candidate BSRs and RPs.

RP paths and SPT paths

[PIM Sparse](#) on page 80 shows two paths for packets from the source for group 239.255.162.1 and a receiver for the group. The source is attached to PIM Sparse device A and the recipient is attached to PIM Sparse device C. PIM Sparse device B is the RP for this multicast group. As a result, the default path for packets from the source to the receiver is through the RP. However, the path through the RP sometimes is not the shortest path. In this case, the shortest path between the source and the receiver is over the direct link between device A and device C, which bypasses the RP (device B).

To optimize PIM traffic, the protocol contains a mechanism for calculating the Shortest Path Tree (SPT) between a given source and receiver. PIM Sparse devices can use the SPT as an alternative to using the RP for forwarding traffic from a source to a receiver. By default, the Extreme device forwards the first packet they receive from a given source to a given receiver using the RP path, but forward subsequent packets from that source to that receiver through the SPT. In [PIM Sparse](#) on page 80, device A forwards the first packet from group 239.255.162.1's source to the destination by sending the packet to device B, which is the RP. Device B then sends the packet to device C. For the second and all future packets that device A receives from the source for the receiver, device A forwards them directly to device C using the SPT path.

Configuring PIM Sparse

To configure a Extreme device for PIM Sparse, perform the following tasks:

- Configure the following global parameter:
 - Enable the PIM Sparse mode of multicast routing.
- Configure the following interface parameters:
 - Configure an IP address on the interface
 - Enable PIM Sparse.
 - Identify the interface as a PIM Sparse border, if applicable.
- Configure the following PIM Sparse global parameters:
 - Identify the Extreme device as a candidate PIM Sparse Bootstrap Router (BSR), if applicable.
 - Identify the Extreme device as a candidate PIM Sparse Rendezvous Point (RP), if applicable.
 - Specify the IP address of the RP (if you want to statically select the RP).

NOTE

It is recommended that you configure the same Extreme device as both the BSR and the RP.

NOTE

It is recommended to use Loopback address as RP instead of VIP for redundancy.

Current limitations

The implementation of PIM Sparse in the current software release has the following limitations:

- PIM Sparse and regular PIM (dense mode) cannot be used on the same interface.

- You cannot configure or display PIM Sparse information using the Web Management Interface. (You can display some general PIM information, but not specific PIM Sparse information.)

Configuring global PIM Sparse parameters

To configure basic global PIM Sparse parameters, enter commands such as the following on each Extreme device within the PIM Sparse domain.

```
device(config)# router pim
```

Syntax: **[no]** router pim

NOTE

You do not need to globally enable IP multicast routing when configuring PIM Sparse.

The command in this example enables IP multicast routing, and enables the PIM Sparse mode of IP multicast routing. The command does not configure the Extreme device as a candidate PIM Sparse Bootstrap Router (BSR) and candidate Rendezvous Point (RP). You can configure a device as a PIM Sparse device without configuring the Extreme device as a candidate BSR and RP. However, if you do configure the device as one of these, it is recommended that you configure the device as both of these. Refer to [Configuring BSRs](#) on page 84.

Entering a **[no]** router pim command does the following:

- Disables PIM.
- Removes all configuration for PIM multicast on a Extreme device (**router pim** level) only.

Enabling PIM Sparse for a specified VRF

To enable PIM for the VRF named "blue", use the following commands.

```
device(config)# router pim vrf blue
```

Syntax: **[no]** router pim [vrf vrf-name]

The **vrf** parameter allows you to configure PIM (PIM-DM and PIM-SM) on the virtual routing instance (VRF) specified by the **vrf-name** variable. All PIM parameters available for the default router instance are configurable for a VRF-based PIM instance.

The **[no]** router pim vrf command behaves in the following manner:

- Entering the **router pim vrf** command to enable PIM does not require a software reload.
- Entering a **no router pim vrf** command removes all configuration for PIM multicast on the specified VRF.

Configuring PIM interface parameters

After you enable IP multicast routing and PIM Sparse at the global level, you must enable it on the individual interfaces connected to the PIM Sparse network.

To enable PIM Sparse mode on an interface, enter commands such as the following.

```
device(config)# interface ethernet 2/2
device(config-if-e10000-2/2)# ip address 207.95.7.1 255.255.255.0
device(config-if-e10000-2/2)# ip pim-sparse
```

Syntax: **[no]** ip pim-sparse

The commands in this example add an IP interface to port 2/2, then enable PIM Sparse on the interface.

If the interface is on the border of the PIM Sparse domain, you also must enter the following command.

```
device(config-if-e10000-2/2)# ip pim border
```

Syntax: [no] ip pim border

Configuring BSRs

In addition to the global and interface parameters described in the previous sections, you need to identify an interface on at least one device as a candidate PIM Sparse Bootstrap router (BSR) and candidate PIM Sparse Rendezvous Point (RP).

NOTE

It is possible to configure the device as only a candidate BSR or RP, but it is recommended that you configure the same interface on the same device as both a BSR and an RP.

This section describes how to configure BSRs. Refer to [Configuring RPs](#) on page 84 for instructions on how to configure RPs.

To configure the device as a candidate BSR, enter commands such as the following.

```
device(config)# router pim
device(config-pim-router)# bsr-candidate ethernet 2/2 30 255
BSR address: 207.95.7.1, hash mask length: 30, priority: 255
```

These commands configure the PIM Sparse interface on port 2/2 as a BSR candidate, with a hash mask length of 30 and a priority of 255. The information shown in italics above is displayed by the CLI after you enter the candidate BSR configuration command.

Syntax: [no] bsr-candidate ethernet slot/portnum | loopback num | tunnel num | ve num hash-mask-length [priority]

The **ethernet slot/portnum | loopback num | ve num** parameters specify the interface. The device will advertise the IP address of the specified interface as a candidate BSR.

- Enter **ethernet slot/portnum** for a physical interface (port).
- Enter **ve num** for a virtual interface.
- Enter **loopback num** for a loopback interface.
- Enter **tunnel num** for a GRE tunnel interface to be configured. The GRE tunnel interface is enabled under the router PIM configuration.

The *hash-mask-length* parameter specifies the number of bits in a group address that are significant when calculating the group-to-RP mapping. You can specify a value from 1-32.

NOTE

it is recommended that you specify 30 for IP version 4 (IPv4) networks.

The *priority* specifies the BSR priority. You can specify a value from 0 - 255. When the election process for BSR takes place, the candidate BSR with the highest priority becomes the BSR. The default is 0.

Configuring RPs

Enter a command such as the following to configure the device as a candidate RP.

```
device(config-pim-router)# rp-candidate ethernet 2/2
```

Syntax: [no] rp-candidate ethernet slot/portnum | loopback num | tunnel num | ve num

The **ethernet slot/portnum | loopback num | ve num** parameters specify the interface. The device will advertise the IP address of the specified interface as a candidate RP.

- Enter **ethernet slot/portnum** for a physical interface (port).

- Enter **ve num** for a virtual interface.
- Enter **loopback num** for a loopback interface.
- Enter **tunnel num** for a GRE tunnel interface to be configured. The GRE tunnel interface is enabled under the router PIM configuration.

By default, this command configures the device as a candidate RP for all group numbers beginning with 224. As a result, the device is a candidate RP for all valid PIM Sparse group numbers. You can change this by adding or deleting specific address ranges. Consider the following when configuring the RP.

- When the candidate RP is configured, before explicitly specifying the groups that it serves, the `c-rp` does, by default, serve all the groups in the PIMSM multicast range, but this includes all groups beginning with **224.x.x.x all the way up to 239.x.x.x**. This is reflected in the "rp-candidate add 224.0.0.0 4" line displayed as part of the runtime configs. This entry will be referred to as the **DEFAULT PREFIX**
- When any group prefix is explicitly added (and the 224.0.0.0/4 prefix itself can also be explicitly added through CLI), the **default prefix** is implicitly removed. Now, the only groups served by the candidate RP, are the groups that have been explicitly added.
- All explicitly added groups can be removed using the "delete" option or "no ... add" option. However, once all the explicitly added groups are deleted from the Candidate RP group prefix list, the **default prefix** becomes active once more. This default group prefix **CANNOT BE REMOVED**.
- It is not possible to punch holes in the group prefix range. For instance executing

```
rp-candidate add 228.0.0.0/16
```

and then,

```
rp-candidate delete 228.0.1.0/24
```

is not permissible. It **cannot** be used to ensure that the rp-candidate will serve all group prefixes in the 228.0.0.0/16 range except those in the 228.0.1.0/24 range.

The following example narrows the group number range for which the device is a candidate RP by explicitly adding a range.

```
device(config-pim-router)# rp-candidate add 224.126.0.0 16
```

Syntax: `[no] rp-candidate add group-addr mask-bits`

The **group-addr mask-bits** specifies the group address and the number of significant bits in the subnet mask. In this example, the device is a candidate RP for all groups that begin with 224.126. When you add a range, you override the default. The device then becomes a candidate RP only for the group address ranges you add.

You also can delete the configured rp-candidate group ranges by entering the following command.

```
device(config-pim-router)# rp-candidate delete 224.126.22.0 24
```

Syntax: `[no] rp-candidate delete group-addr mask-bits`

The usage of the **group-addr mask-bits** parameter is the same as for the **rp-candidate add** command.

Updating PIM-Sparse forwarding entries with new RP configuration

If you make changes to your static RP configuration, the entries in the PIM-Sparse multicast forwarding table continue to use the old RP configuration until they are aged out.

The **clear pim rp-map** command allows you to update the entries in the static multicast forwarding table immediately after making RP configuration changes. This command is meant to be used with **rp-address** command.

To update the entries in a PIM sparse static multicast forwarding table with new RP configuration, enter the following command at the privileged EXEC level of the CLI.

```
device(config)# clear ip pim rp-map
```

Syntax: `clear ip pim [vrf vrf-name] rp-map`

Use the **vrf** option to clear the PIM sparse static multicast forwarding table for a VRF instance specified by the *vrf-name* variable.

Statically specifying the RP

It is recommended that you use the PIM Sparse protocol's RP election process so that a backup RP can automatically take over if the active RP router becomes unavailable. However, if you do not want the RP to be selected by the RP election process but instead you want to explicitly identify the RP by P address, use the **rp-address** command.

If you explicitly specify the RP, the device uses the specified RP for all group-to-RP mappings and overrides the set of candidate RPs supplied by the BSR.

NOTE

Specify the same IP address as the RP on all PIM Sparse devices within the PIM Sparse domain. Make sure the device is on the backbone or is otherwise well connected to the rest of the network.

To specify the IP address of the RP, enter commands such as the following.

```
device(config)# router pim
device(config-pim-router)# rp-address 207.95.7.1
```

Syntax: `[no] rp-address ip-addr`

The **ip-addr** parameter specifies the IP address of the RP.

The command in this example identifies the device interface at IP address 207.95.7.1 as the RP for the PIM Sparse domain. The device uses the specified RP and ignore group-to-RP mappings received from the BSR.

ACL based RP assignment

The **rp-address** command allows multiple static RP configurations. For each static RP, an ACL can be given as an option to define the multicast address ranges that the static RP permit or deny to serve.

A static RP by default serves the range of 224.0.0.0/4 if the RP is configured without an ACL name. If an ACL name is given but the ACL is not defined, the static RP is set to inactive mode and it will not cover any multicast group ranges.

The optional static RP ACL can be configured as a standard ACL or as an extended ACL. For an extended ACL, the destination filter will be used to derive the multicast group range and all other filters are ignored. The content of the ACL needs to be defined in the order of prefix length; the longest prefix must be placed at the top of the ACL definition.

If there are overlapping group ranges among the static RPs, the static RP with the longest prefix match is selected. If more than one static RP covers the exact same group range, the highest IP static RP will be used.

Configuration considerations:

- The Static RP has higher precedence over RP learnt from the BSR.
- There is a limit of 64 static RPs in the systems.

Configuring an ACL based RP assignment

To configure an ACL based RP assignment, enter commands such as the following.

```
device(config)# router pim
device(config-pim-router)# rp-address 130.1.1.1 acl1
```

Syntax: [no] **rp-address** *ip_address* [*acl_name_or_id*] [*vrf vrf-name*]

Use the **ip address** parameter to specify the IP address of the device you want to designate as an RP device.

Use the **acl name** or **id** (optional) parameter to specify the name or ID of the ACL that specifies which multicast groups use this RP.

Use the **vrf** parameter to specify a VRF instance.

Displaying the static RP

Use the **show ip pim rp-set** command to display static RP and the associated group ranges.

```
device(config)# show ip pim rp-set
Static RP and associated group ranges
-----
Static RP count: 4
130.1.1.1
120.1.1.1
120.2.1.1
124.1.1.1
Number of group prefixes Learnt from BSR: 0
No RP-Set present.
```

Use the **show ip pim rp-map** command to display all current multicast group addresses to RP address mapping.

```
device(config)# show ip pim rp-map
Number of group-to-RP mappings: 5
  Group address      RP address
-----
1    230.0.0.1         100.1.1.1
2    230.0.0.2         100.1.1.1
3    230.0.0.3         100.1.1.1
4    230.0.0.4         100.1.1.1
5    230.0.0.5         100.1.1.1
```

Route selection precedence for multicast

The **route-precedence** command lets you specify a precedence table that dictates how routes are selected for multicast.

Configuration considerations:

PIM must be enabled at the global level.

Configuring route precedence by specifying route types

The **route precedence** command lets you control the selection of routes based on the following route types:

- Non-default route from the mRTM
- Default route from the mRTM
- Non-default route from the uRTM
- Default route from the uRTM

Use this command to specify an option for all of the precedence levels.

To specify a non-default route from the mRTM, then a non-default route from the uRTM, then a default route from the mRTM, and then a default route from the uRTM, enter commands such as the following.

```
device(config)# router pim
device(config-pim-router)# route-precedence mc-non-default uc-non-default mc-default uc-default
```

The **none** option may be used to fill up the precedence table in order to ignore certain types of routes. To use the unicast default route for multicast, enter commands such as the following.

```
device(config)# router pim
device(config-pim-router)# route-precedence mc-non-default mc-default uc-non-default none
```

Syntax: [no] route-precedence [mc-non-default | mc-default | uc-non-default | uc-default | none]

Default value: route-precedence mc-non-default mc-default uc-non-default uc-default

Use the **mc-non-default** parameter to specify a multicast non-default route.

Use the **mc-default** parameter to specify a multicast default route.

Use the **uc-non-default** parameter to specify a unicast non-default route.

Use the **uc-default** parameter to specify a unicast default route.

Use the **none** parameter to ignore certain types of routes.

The **no** form of this command removes the configuration.

Displaying the route selection

Use the **show ip pim sparse** command to display the current route selection. This example shows the default route precedence selection.

```
device(config)# show ip pim sparse
Global PIM Sparse Mode Settings
Maximum Mcache      : No limit      Current Count       : 1400
Hello interval      : 1              Neighbor timeout    : 3
Join/Prune interval : 60             Inactivity interval : 180
Hardware Drop Enabled : Yes          Prune Wait Interval : 3
Bootstrap Msg interval : 60         Candidate-RP Msg interval : 60
Register Suppress Time : 60         Register Probe Time : 10
Register Stop Delay   : 60             Register Suppress interval : 60
SSM Enabled          : No              SPT Threshold       : Infinity
Route Precedence     : mc-non-default mc-default uc-non-default uc-default
```

In this example, the route precedence selection is multicast non-default, then unicast non-default, then multicast default, and then unicast default.

NOTE

In the Brocade implementation, the LHR directly sends the PIM joins towards RP instead of DR in a broadcast network. You are advised to enable PIM in all the routers in the broadcast domain because the Unicast RTM currently only stores the BEST METRIC path towards a particular IP address and that is all that Multicast has access to when it polls the RTM. If the next hop neighbor for that path is NOT PIM enabled, then PIM will not know where to pick up the next available path (since the RTM does not hold it).

Changing the Shortest Path Tree (SPT) threshold

In a typical PIM Sparse domain, there may be two or more paths from a DR (designated router) for a multicast source to a PIM group receiver:

- **Path through the RP** - This is the path the device uses the first time it receives traffic for a PIM group. However, the path through the RP may not be the shortest path from the device to the receiver.
- **Shortest Path** - Each PIM Sparse device that is a DR for a multicast source calculates a shortest path tree (SPT) to all the PIM Sparse group receivers within the domain, with the device itself as the root of the tree. The first time a device configured as a PIM router receives a packet for a PIM receiver, the device sends the packet to the RP for the group. The device also calculates the SPT from itself to the receiver. The next time the device receives a PIM Sparse packet for the receiver, the device sends the packet toward the receiver using the shortest route, which may not pass through the RP.

By default, the device switches from the RP to the SPT after receiving the first packet for a given PIM Sparse group. The device maintains a separate counter for each PIM Sparse source-group pair.

After the device receives a packet for a given source-group pair, it starts a PIM data timer for that source-group pair. If the device does not receive another packet for the source-group pair before the timer expires, it reverts to using the RP for the next packet received for the source-group pair. In accordance with the PIM Sparse RFC recommendation, the timer is 210 seconds and is not configurable. The counter is reset to zero each time the device receives a packet for the source-group pair.

You can change the number of packets that the device sends using the RP before switching to using the SPT by entering commands such as the following.

```
device(config)# router pim
device(config-pim-router)# spt-threshold 1000
```

Syntax: [no] spt-threshold infinity | num

The **infinity** | **num** parameter specifies the number of packets. If you specify **infinity**, the device sends packets using the RP indefinitely and does not switch over to the SPT. If you enter a specific number of packets, the device does not switch over to using the SPT until it has sent the number of packets you specify using the RP.

Configuring PIM-SM (*,g) forwarding

By default, the device supports only (s,g) hardware routing. This is adequate in network topologies where there a limited number of multicast flows, and most SPT and RPF paths diverge from the PIM Last Hop (PIM LH) from the PIM First Hop (PIM FH).

If however, the RPF path from a PIM LH to the PIM RP and the source is the same, the intermediate devices between RP and LH can be optimized to aggregate multicast flows destined to the same group address to use a single (*,g) CAM entry, instead of consuming an (s,g) hardware entry for each flow. This reduces CAM usage as well as the number of IPCs between the interface and management modules to manage the individual (s,g) states.

For example, where a service provider is provisioning multicast service, with only the route of PIM RP being visible to customers, and routes to the sources are all default, the (*,g) hardware entry can help optimize system resources by keeping the traffic on the shared tree.

NOTE

It is recommended to configure the **spt-threshold infinity** command beginning from the PIM LH router, then to the intermediate PIM routers and finally to the PIM RP.

Configuration details

To enable this feature, you must explicitly configure the **spt-threshold infinity** command on all multicast nodes, as shown in the following example.

```
device(config)# router pim
device(config-pim-router)# spt-threshold infinity
```

Syntax: [no] spt-threshold infinity

This configuration option forces PIM-SM to distribute multicast traffic on the share tree only. PIM devices from RP to LH maintain only (*,g) states, regardless of the source location in the network topology. PIM LH will not initiate (s,g) SPT switchover, so there are no (s,g) joins generated from LH. Devices in the share tree, eg, devices downstream from RP, create a (*,g) hardware forwarding state to switch multicast flows for the group.

Changing the PIM Join and Prune message interval

By default, the device sends PIM Sparse Join or Prune messages every 60 seconds. These messages inform other PIM Sparse devices about clients who want to become receivers (Join) or stop being receivers (Prune) for PIM Sparse groups.

NOTE

Use the same Join or Prune message interval on all the PIM Sparse devices in the PIM Sparse domain. If the devices do not all use the same timer interval, the performance of PIM Sparse can be adversely affected.

To change the Join or Prune interval, enter commands such as the following.

```
device(config)# router pim
device(config-pim-router)# message-interval 30
```

Syntax: [no] message-interval num

The **num** parameter specifies the number of seconds and can from 10 - 65535. The default is 60.

PIM multinet

Brocade devices support PIM over secondary addresses in IPv4.

Whenever a secondary address is configured on a interface, all the secondary addresses configured on the interface are sent out on the PIM Hello using the secondary address option.

Whenever a receiver uses a secondary address as its source and sends a IGMP group report, the PIM join and prunes are propagated up the network.

Whenever a secondary address is configured as a RP, the packets are processed appropriately

Displaying the secondary address

In this example the PIM neighbor on e 1/11 has multiple IP addresses configured on the interface.

```
device#show ip pim neighbor
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Port  |PhyPort |Neighbor |Holdtime| T |PropDelay|Override|Age|UpTime  | VRF  |Prio
|      |         |         | sec    | Bit| msec    | msec   |sec|         |      |
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
e1/11 e1/11   20.1.1.1 105     1  500     3000   18 00:00:50 default 1
      + 30.1.1.1
e1/12 e1/12   192.168.2.1 105    1  500     3000   18 2d 02:34:19 default 1
e1/15 e1/15   192.168.7.3 105    1  500     3000   7 2d 23:53:07 default 1
On the PIM interface when multiple IP addresses are configured, the lowest IP address is
```

chosen as the primary address.
That is the reason that ethernet 1/11 has 20.1.1.1 as the address of the neighbor address.

Syntax: show ip pim neighbor

Refer to [Displaying multicast neighbor information](#) on page 98 for output descriptions.

Multicast Outgoing Interface (OIF) list optimization

Each multicast route entry maintains a list of outgoing interfaces (OIF List) to which an incoming multicast data packet matching the route is replicated. In hardware-forwarded route entries, these OIF lists are stored inside the hardware in replication tables which are limited in size. In many deployment scenarios, more than one multicast route can have identical OIF lists and can optimize usage of the replication table entries by sharing them across multiple multicast routes. Multicast OIF list optimization keeps track of all the OIF lists in the system. It manages the hardware replication resources optimally, in real time, by dynamically assigning or re-assigning resources to multicast route entries to suit their current OIF list requirements, while maximizing resource sharing.

Displaying PIM Sparse configuration information and statistics

You can display the following PIM Sparse information:

- Basic PIM Sparse configuration information
- Group information
- BSR information
- Candidate RP information
- RP-to-group mappings
- RP information for a PIM Sparse group
- RP set list
- PIM neighbor information
- The PIM flow cache
- The PIM multicast cache
- PIM traffic statistics
- PIM counter statistics

Displaying basic PIM Sparse configuration information

To display PIM Sparse configuration information, enter the following command at any CLI level.

```
device(config)# show ip pim sparse
Global PIM Sparse Mode Settings
  Maximum Mcache          : 0           Current Count          : 0
  Hello interval          : 30           Neighbor timeout       : 105
  Join/Prune interval     : 60           Inactivity interval   : 180
  Hardware Drop Enabled   : Yes          Prune Wait Interval   : 3
  Bootstrap Msg interval : 60           Candidate-RP Msg interval : 60
  Register Suppress Time  : 60           Register Probe Time   : 10
  Register Stop Delay     : 60           Register Suppress interval : 60
  SSM Enabled             : No            SPT Threshold         : 1
  Route Precedence        : mc-non-default mc-default uc-non-default uc-default
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
Interface|Local   |Ver |St | Designated Router |TTL|Multicast| VRF  | DR |Override
         |Address |   |   | Address           |Thr|Boundary |     |   |Prio|Interval
```

```

-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
v3      3.1.1.2   SMv2 Ena Itself          1 None default 1 3000ms
v4      4.1.1.2   SMv2 Ena Itself          1 None default 1 3000ms
v10     10.1.1.1  SMv2 Ena Itself          1 None default 1 3000ms
v12     12.1.1.1  SMv2 Ena Itself          1 None default 1 3000ms
v33     33.1.1.1  SMv2 Ena 33.1.1.2      1/14 1 None default 1 3000ms
11      1.1.1.3   SMv2 Ena Itself          1 None default 1 3000ms
Total Number of Interfaces : 6

```

Syntax: show ip pim [vrf vrf-name] sparse

The **vrf** option allows you to display PIM sparse configuration information for the VRF instance identified by the *vrf-name* variable.

This example shows the PIM Sparse configuration information on PIM Sparse device A in [PIM Sparse](#) on page 80.

The table below shows the information displayed by the **show ip pim sparse** command.

TABLE 10 Output of the show ip pim sparse command

This field...	Displays...
Global PIM Sparse mode settings	
Hello interval	How frequently the device sends PIM Sparse hello messages to its PIM Sparse neighbors. This field also shows the number of seconds between hello messages. PIM Sparse devices use hello messages to discover each another.
Neighbor timeout	How many seconds the device waits for a hello message from a neighbor before determining that the neighbor is no longer present and removing cached PIM Sparse forwarding entries for the neighbor.
Bootstrap Msg interval	How frequently the BSR configured on the device sends the RP set to the RPs within the PIM Sparse domain. The RP set is a list of candidate RPs and their group prefixes. A candidate RP group prefix indicates the range of PIM Sparse group numbers for which it can be an RP. NOTE This field contains a value only if an interface on the device is elected to be the BSR. Otherwise, the field is blank.
Candidate-RP Advertisement interval	How frequently the candidate PR configured on the device sends candidate RP advertisement messages to the BSR. NOTE This field contains a value only if an interface on the device is configured as a candidate RP. Otherwise, the field is blank.
Join or Prune interval	How frequently the device sends PIM Sparse Join or Prune messages for the multicast groups it is forwarding. This field also shows the number of seconds between Join or Prune messages. The device sends Join or Prune messages on behalf of multicast receivers who want to join or leave a PIM Sparse group. When forwarding packets from PIM Sparse sources, the device sends the packets only on the interfaces on which it has received join requests in Join or Prune messages for the source group. You can change the Join or Prune interval. Refer to Changing the PIM Join and Prune message interval on page 90.
SPT Threshold	The number of packets the device sends using the path through the RP before switching to using the SPT path.
Inactivity Interval	
SSM Enabled	If yes, source-specific multicast is configured globally on this device.
SSM Group Range	The SSM range of IP multicast addresses. By default this range will be 232/8 as assigned by the Internet Assigned Numbers Authority (IANA) for use with SSM. Other values can be configured.
PIM Sparse interface information	

TABLE 10 Output of the show ip pim sparse command (continued)

This field...	Displays...
<p>NOTE You also can display IP multicast interface information using the show ip pim interface command. However, this command lists all IP multicast interfaces, including regular PIM (dense mode) interface. The show ip pim sparse command lists only the PIM Sparse interfaces.</p>	
Interface	<p>The type of interface and the interface number. The interface type can be one of the following:</p> <ul style="list-style-type: none"> Ethernet VE <p>The number is either a port number (and slot number if applicable) or the virtual interface (VE) number.</p>
TTL Threshold	<p>Following the TTL threshold value, the interface state is listed. The interface state can be one of the following:</p> <ul style="list-style-type: none"> Disabled Enabled
Local Address	Indicates the IP address configured on the port or virtual interface.
Mode	
Version	
Designated Router	

Displaying a list of multicast groups

To display PIM group information, enter the following command at any CLI level.

```
device(config)# show ip pim group
Total number of groups for VRF default-vrf: 7
1   Group 226.0.34.0
    Group member at e2/9: v59
    Group member at e1/16: v57
2   Group 226.0.77.0
    Group member at e2/9: v59
    Group member at e1/16: v57
3   Group 226.0.120.0
    Group member at e2/9: v59
    Group member at e1/16: v57
4   Group 226.0.163.0
    Group member at e2/9: v59
    Group member at e1/16: v57
5   Group 226.0.206.0
    Group member at e2/9: v59
    Group member at e1/16: v57
6   Group 226.0.249.0
    Group member at e2/9: v59
    Group member at e1/16: v57
7   Group 226.0.30.0
    Group member at e2/9: v59
    Group member at e1/16: v57
device(config)#
```

Syntax: **show ip pim [vrf vrf-name] group**

The **vrf** option allows you to display PIM group information for the VRF instance identified by the **vrf-name** variable.

Table 11 describes the output from this command.

TABLE 11 Output from the `show ip pim vrf group` command

This field...	Displays...
Total number of Groups	Lists the total number of IP multicast groups the device is forwarding. NOTE This list can include groups that are not PIM Sparse groups. If interfaces on the device are configured for regular PIM (dense mode), these groups are listed too.
Index	The index number of the table entry in the display.
Group	The multicast group address
Ports	The device ports connected to the receivers of the groups.

Displaying BSR information

To display BSR information, enter the following command at any CLI level.

```
device(config)# show ip pim bsr
PIMv2 Bootstrap information for Vrf Instance : default-vrf
-----
This system is the Elected BSR
BSR address: 1.51.51.1. Hash Mask Length 32. Priority 255.
Next bootstrap message in 00:01:00
Configuration:
Candidate loopback 2 (Address 1.51.51.1). Hash Mask Length 32. Priority 255.
Next Candidate-RP-advertisement in 00:01:00
RP: 1.51.51.1
group prefixes:
224.0.0.0 / 4
Candidate-RP-advertisement period: 60
c(config)#
```

This example shows information displayed on a device that has been elected as the BSR. The next example shows information displayed on a device that is not the BSR. Notice that some fields shown in the example above do not appear in the example below.

```
device(config)#show ip pim bsr
PIMv2 Bootstrap information for Vrf Instance : default-vrf
-----
BSR address: 1.51.51.1. Hash Mask Length 32. Priority 255.
Next Candidate-RP-advertisement in 00:00:30
RP: 1.51.51.3
group prefixes:
224.0.0.0 / 4
Candidate-RP-advertisement period: 60
device(config)#
```

Syntax: `show ip pim [vrf vrf-name] bsr`

The `vrf` option allows you to display BSR information for the VRF instance identified by the `vrf-name` variable.

[Table 12](#) describes the output from this command.

TABLE 12 Output from the `show ip pim bsr` command

This field...	Displays...
BSR address	The IP address of the interface configured as the PIM Sparse Bootstrap Router (BSR).
Uptime	The amount of time the BSR has been running. NOTE This field appears only if this device is the BSR.

TABLE 12 Output from the `show ip pim bsr` command (continued)

This field...	Displays...
BSR priority	The priority assigned to the interface for use during the BSR election process. During BSR election, the priorities of the candidate BSRs are compared and the interface with the highest BSR priority becomes the BSR.
Hash mask length	The number of significant bits in the IP multicast group comparison mask. This mask determines the IP multicast group numbers for which the device can be a BSR. The default is 32 bits, which allows the device to be a BSR for any valid IP multicast group number. NOTE This field appears only if this device is a candidate BSR.
Next bootstrap message in	Indicates how many seconds will pass before the BSR sends the next bootstrap message. NOTE This field appears only if this device is the BSR.
Next Candidate-RP-advertisement message in	Indicates how many seconds will pass before the BSR sends the next candidate RP advertisement message. NOTE This field appears only if this device is a candidate BSR.
RP	Indicates the IP address of the Rendezvous Point (RP). NOTE This field appears only if this device is a candidate BSR.
group prefixes	Indicates the multicast groups for which the RP listed by the previous field is a candidate RP. NOTE This field appears only if this device is a candidate BSR.
Candidate-RP-advertisement period	Indicates how frequently the BSR sends candidate RP advertisement messages. NOTE This field appears only if this device is a candidate BSR.

Displaying candidate RP information

To display candidate RP information, enter the following command at any CLI level.

```
device# show ip pim rp-candidate
Next Candidate-RP-advertisement in 00:00:10
  RP: 207.95.7.1
    group prefixes:
      224.0.0.0 / 4
Candidate-RP-advertisement period: 60
```

This example shows information displayed on a device that is a candidate RP. The next example shows the message displayed on a device that is not a candidate RP.

```
device# show ip pim rp-candidate
```

This system is not a Candidate-RP.

Syntax: `show ip pim rp-candidate vrf-name`

This command displays candidate RP information for the VRF instance identified by the **vrf-name** variable.

Table 13 describes the output from this command.

TABLE 13 Output from the `show ip pim rp-candidate` command

This field...	Displays...
Candidate-RP-advertisement in	Indicates how many seconds will pass before the BSR sends the next RP message. NOTE This field appears only if this device is a candidate RP.
RP	Indicates the IP address of the Rendezvous Point (RP). NOTE This field appears only if this device is a candidate RP.
group prefixes	Indicates the multicast groups for which the RP listed by the previous field is a candidate RP. NOTE This field appears only if this device is a candidate RP.
Candidate-RP-advertisement period	Indicates how frequently the BSR sends candidate RP advertisement messages. NOTE This field appears only if this device is a candidate RP.

Displaying RP-to-group mappings

To display RP-to-group-mappings, enter the following command at any CLI level.

```
device# show ip pim rp-map
Number of group-to-RP mappings: 6
Group address      RP address
-----
 1 239.255.163.1   99.99.99.5
 2 239.255.163.2   99.99.99.5
 3 239.255.163.3   99.99.99.5
 4 239.255.162.1   99.99.99.5
 5 239.255.162.2   43.43.43.1
 6 239.255.162.3   99.99.99.5
```

Syntax: `show ip pim [vrf vrf-name] rp-map`

The **vrf** option allows you to display candidate RP-to-group mappings for the VRF instance identified by the **vrf-name** variable.

This display shows the following information.

TABLE 14 Output of the `show ip pim rp-map` command

This field...	Displays...
Group address	Indicates the PIM Sparse multicast group address using the listed RP.
RP address	Indicates the IP address of the Rendezvous Point (RP) for the listed PIM Sparse group.

Displaying RP Information for a PIM Sparse group

To display RP information for a PIM Sparse group, enter the following command at any CLI level.

```
device# show ip pim rp-hash 239.255.162.1
RP: 207.95.7.1, v2
Info source: 207.95.7.1, via bootstrap
```

Syntax: `show ip pim [vrf vrf-name] rp-hash group-addr`

The `vrf` option allows you to display RP information for the VRF instance identified by the `vrf-name` variable.

The `group-addr` parameter is the address of a PIM Sparse IP multicast group.

Table 15 describes the output from this command.

TABLE 15 Output from the `show ip pim` command

This field...	Displays...
RP	Indicates the IP address of the Rendezvous Point (RP) for the specified PIM Sparse group, followed by the port or virtual interface through which this device learned the identity of the RP.
Info source	Indicates the IP address on which the RP information was received, followed by the IP address through which this device learned the identity of the RP.

Displaying the RP set list

To display the RP set list, enter the following command at any CLI level.

```
device(config)# show ip pim rp-set
Static RP
-----
Static RP count: 2
1.51.51.4
1.51.51.5
Number of group prefixes Learnt from BSR: 1
Group prefix = 224.0.0.0/4      # RPs: 2
  RP 1: 1.51.51.1  priority=0  age=60  holdtime=150
  RP 2: 1.51.51.3  priority=0  age=30  holdtime=150
device(config)#
```

Syntax: `show ip pim [vrf vrf-name] rp-set`

The `vrf` option allows you to display the RP set list for the VRF instance identified by the `vrf-name` variable.

Table 16 describes the output from this command.

TABLE 16 Output from the `show ip pim vrf rp-set` command

This field...	Displays...
Number of group prefixes	The number of PIM Sparse group prefixes for which the RP is responsible.
Group prefix	Indicates the multicast groups for which the RP listed by the previous field is a candidate RP.

TABLE 16 Output from the **show ip pim vrf rp-set** command (continued)

This field...	Displays...
RP's expected or received	Indicates how many RPs were expected and received in the latest bootstrap message.
RP num	Indicates the RP number. If there are multiple RPs in the PIM Sparse domain, a line of information for each RP is listed, in ascending numerical order.
priority	The RP priority of the candidate RP. During the election process, the candidate RP with the highest priority is elected as the RP.
age	The age (in seconds) of this RP-set. NOTE If this device is not a BSR, this field contains zero. Only the BSR ages the RP-set.

Displaying multicast neighbor information

To display information about PIM neighbors, enter the following command at any CLI level.

```
device(config)# show ip pim nbr
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
Port | PhyPort | Neighbor | Holdtime | T | PropDelay | Override | Age | UpTime | VRF | Prio
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
v2   | e1/1    | 2.1.1.2 | 105      | 1 | 500       | 3000    | 0   | 00:44:10 | default-vrf | 1
v4   | e2/2    | 4.1.1.2 | 105      | 1 | 500       | 3000    | 10  | 00:42:50 | default-vrf | 1
v5   | e1/4    | 5.1.1.2 | 105      | 1 | 500       | 3000    | 0   | 00:44:00 | default-vrf | 1
v22  | e1/1    | 22.1.1.1 | 105      | 1 | 500       | 3000    | 0   | 00:44:10 | default-vrf | 1
Total Number of Neighbors : 4
device(config)#
```

Syntax: **show ip pim [vrf *vrf-name*] neighbor**

The **vrf** option allows you to display information about the PIM neighbors for the VRF instance identified by the **vrf-name** variable.

Table 17 describes the output from this command.

TABLE 17 Output from the **show ip pim vrf neighbor** command

This field...	Displays...
Port	The interface through which the device is connected to the neighbor.
Neighbor	The IP interface of the PIM neighbor.
Holdtime sec	Indicates how many seconds the neighbor wants this device to hold the entry for this neighbor in memory. The neighbor sends the Hold Time in Hello packets: <ul style="list-style-type: none"> If the device receives a new Hello packet before the Hold Time received in the previous packet expires, the device updates its table entry for the neighbor. If the device does not receive a new Hello packet from the neighbor before the Hold time expires, the device assumes the neighbor is no longer available and removes the entry for the neighbor.
Age sec	The number of seconds since the device received the last hello message from the neighbor.
UpTime sec	The number of seconds the PIM neighbor has been up. This timer starts when the device receives the first Hello messages from the neighbor.

TABLE 17 Output from the `show ip pim vrf neighbor` command (continued)

This field...	Displays...
VRF	The VRF in which the interface is configured. This can be a VRF that the port was assigned to or the default VRF of the device.
Priority	The DR priority that is used in the DR election process. This can be a configured value or the default value of 1.

Displaying the PIM multicast cache

To display the PIM multicast cache, enter the following command at any CLI level.

```
device(config)# show ip pim mcache 54.1.1.10 226.0.1.0
IP Multicast Mcache Table
Entry Flags      : SM - Sparse Mode, SSM - Source Specific Mutlicast, DM - Dense Mode
                  RPT - RPT Bit, SPT - SPT Bit, LSRC - Local Source, LRCV - Local Receiver
                  HW - HW Forwarding Enabled, FAST - Resource Allocated, TAG - Need For Replication Entry
                  REGPROB - Register In Progress, REGSUPP - Register Suppression Timer
                  MSDPADV - Advertise MSDP, NEEDRTE - Route Required for Src/RP, PRUN - DM Prune Upstream
Interface Flags: IM - Immediate, IH - Inherited
                  MJ - Membership Join, MI - Membership Include, ME - Membership Exclude
                  BR - Blocked RPT, BA - Blocked Assert, BF - Blocked Filter, BI - Blocked IIF
1 (54.1.1.10, 226.0.1.0) in v52 (e2/6), Uptime 00:45:55, Rate 115 (SM)
upstream neighbor 52.1.1.1
Flags (0xf006c4e1) SM SPT LRCV HW FAST TAG MSDPADV
fast ports: ethe 1/16 ethe 2/9
AgeSltMsk: 00000002, FID: 0x8225, MVID: 257, RegPkt: 44, AvgRate: 114, profile: none
Forwarding_oif: 2, Immediate_oif: 1, Blocked_oif: 2
L3 (HW) 2:
  e1/16(VL57), 00:43:17/0, Flags: MJ
  e2/9(VL59), 00:44:45/168, Flags: IM IH MJ
Blocked OIF 2:
  e1/2(VL53), 00:44:00/209, Flags: IH BR
  e2/6(VL52), 00:44:01/209, Flags: IH BR BI
device(config)#
```

Syntax: `show ip pim [vrf vrf-name] mcache`

The `vrf` option allows you to display the PIM multicast cache for the VRF instance identified by the `vrf-name` variable.

TABLE 18 Output fields from the `show ip pim mcache` command

Field	Description
Total entries in mcache	Shows the total number of PIM mcache entries
MJ	Membership Join
MI	Membership Include
ME	Membership Exclude - Legend for the mcache entry printed once per page, it gives the explanation of each of the flags used in the entry.
BR	Blocked RPT
BA	Blocked Assert
BF	Blocked Filter
BI	Blocked IIF
Uptime	Shows the entry uptime
Rate	Shows the Rate at which packets are ingressing for this entry
upstream neighbor	Shows the upstream neighbor for the Source/RP based on the type of entry. For (*,G) it shows the upstream neighbor towards the RP. For (S,G) entries it shows the upstream neighbor towards the source.

TABLE 18 Output fields from the **show ip pim mcache** command (continued)

Field	Description
Flags	<p>Flags Represent Entry flags in hex format in the braces. And indicates the meaning of the flags set in abbreviated string whose explanations are as below. Only shows the flags which are set.</p> <p>SM - Shows If the entry is created by PIM Sparse Mode</p> <p>DM - Shows If DM mode entry is enabled</p> <p>SSM - Shows If the SSM mode entry is enabled</p> <p>RPT - Shows If the entry is on the Rendezvous Point (RP)</p> <p>SPT - Shows If the entry is on the source tree</p> <p>LSRC - Shows If the source is in a directly-connected interface</p> <p>LRcv - Shows If the receiver is directly connected to the router</p> <p>REG - if the data registration is in progress</p> <p>L2REG - if the source is directly connected to the router</p> <p>REGSUPP - if the register suppression timer is running</p> <p>RegProbe</p> <p>HW - Shows If the candidate for hardware forwarding is enabled</p> <p>FAST - Shows If the resources are allocated for hardware forwarding</p> <p>TAG - Shows If there is a need for allocating entries from the replication table</p> <p>MSDPADV - Shows If RP is responsible for the source and must be advertised to its peers.</p> <p>NEEDRTE - Shows If there is no route to the source and RP is available</p> <p>PRUNE - Shows If PIM DM Prune to upstream is required</p>
RP	Show the IP address of the RP.
fast ports	Shows forwarding port mask.
AgeSlitMsk	Shows the slot number on which MP expects ingress traffic.
FID, MVID	Shows the resources that are allocated for forwarding the entry.
RegPkt	Shows Count of Packets forwarded due to the Register decapsulation.
AvgRate	Shows the average Rate of packets ingressing for this entry over 30 seconds.
Profile	Shows the Profile ID associated with the Stream.
Number of matching entries	Shows the total number of mcache entries matching a particular multicast filter specified.
Outgoing interfaces Section	<p>This section consists of three parts. L3 OIFs, L2OIFs and Blocked OIFs. And each section has Format of L3/L2/Blocked followed by (HW/SW) followed by count of the number of OIF in each section.</p> <p>Additionally, each section displays the OIFs one per line. And shows the OIF in the format eth/Tr(Vlan) followed by uptime/expiry time, followed by the Flags associated with each OIF.</p>
L3	Show whether the traffic is routed out of the interface.
L2	Show whether the traffic is switched out of the interface.
HW	Show whether the entry is hardware forwarded.
SW	Show whether the entry is software forwarded

TABLE 18 Output fields from the **show ip pim mcache** command (continued)

Field	Description
Eth/Tr(VL1)	Shows the outgoing interface on the specified VLAN.
Flags (explanation of flags in the OIF section)	Shows the flags set in each of the Outgoing interface in abbreviated string format whose explanations are as below. Legend of this shown at the top of each entry IM - Immediate IH - Inherited MJ - Membership Join MI - Membership Include ME - Membership Exclude BR - Blocked due to SG RPT BA - Blocked due to Assert BF - Blocked due to Filter BI - Blocked IIF (Incoming interface) matches OIF

Displaying the PIM multicast cache for MVID

To display the PIM multicast cache for a specified mvid, enter the following command at any CLI level.

```

device# show ip pim mcache mvid 1
Total entries in mcache: 17
 1 (49.1.1.100, 229.0.0.1) in v49 (tag e4/9), Uptime 00:02:36 Rate 0
   Sparse Mode, RPT=0 SPT=1 Reg=0 L2Reg=1 RegSupp=0 RegProbe=0 LSrc=1 LRcv=1
   Source is directly connected. RP 192.168.1.1
   num_oifs = 1
   immediate_oifs = 0, inherited_oifs = 1, blocked_oifs = 0
     52,4/12(00:02:36/0) Flags:00000004
   fast ports ethe 4/12
   L3 (HW) 1: e4/12(VL52)
   Flags (0x7046c8e1)
     sm=1 ssm=0 hw=1 fast=1 slow=0 leaf=0 prun=0 tag=1 needRte=0 msdp_adv=1
   AgeSltMsk=00000008, FID: 0x800e MVID: 1, RegPkt 0, AvgRate 0 profile: none
 2 (49.1.1.108, 229.0.0.1) in v49 (tag e4/9), Uptime 00:03:51 Rate 11257
   Sparse Mode, RPT=0 SPT=1 Reg=0 L2Reg=1 RegSupp=0 RegProbe=0 LSrc=1 LRcv=1
   Source is directly connected. RP 192.168.1.1
   num_oifs = 1
   immediate_oifs = 0, inherited_oifs = 1, blocked_oifs = 0
     52,4/12(00:03:51/0) Flags:00000004
   fast ports ethe 4/12
   L3 (HW) 1: e4/12(VL52)
   Flags (0x7046c8e1)
     sm=1 ssm=0 hw=1 fast=1 slow=0 leaf=0 prun=0 tag=1 needRte=0 msdp_adv=1
   AgeSltMsk=00000008, FID: 0x800e MVID: 1, RegPkt 0, AvgRate 11257 profile: none
 3 (49.1.1.108, 229.0.0.2) in v49 (tag e4/9), Uptime 00:03:51 Rate 11257
   Sparse Mode, RPT=0 SPT=1 Reg=0 L2Reg=1 RegSupp=0 RegProbe=0 LSrc=1 LRcv=1
   Source is directly connected. RP 192.168.1.1
   num_oifs = 1
   immediate_oifs = 0, inherited_oifs = 1, blocked_oifs = 0
     52,4/12(00:03:51/0) Flags:00000004
   fast ports ethe 4/12
   L3 (HW) 1: e4/12(VL52)
   Flags (0x7046c8e1)
     sm=1 ssm=0 hw=1 fast=1 slow=0 leaf=0 prun=0 tag=1 needRte=0 msdp_adv=1
   AgeSltMsk=00000008, FID: 0x800e MVID: 1, RegPkt 0, AvgRate 11257 profile: none
Number of matching entries: 3

```

Syntax: `show ip pim mcache mvid mvid`

The *mvid* variable allows you to display an entry that matches a specified mvid.

Table 19 describes the output parameters of the **show ip pim mcache mvid 1** command.

TABLE 19 Output parameters of the **show ip pim mcache mvid 1** command

Field	Description
Total entries in mcache	Shows the total number of PIM mcache entries.
v	Shows the interface name.
tag e	Shows the VLAN tagged ethernet interface.
Uptime	Shows the entry uptime.
Rate	Shows the total packet count.
Sparse mode	Shows that the PIM sparse mode is enabled.
RPT	Sets the flag to 1, if the entry is on the RP tree, else sets the flag to 0.
SPT	Sets the flag to 1, if the entry is on the source tree, or else sets the flag to 0.
Reg	Sets the flag to 1, if the data registration is in progress, or else sets the flag to 0.
L2Reg	Sets the flag to 1, if the source is directly connected to the router, or else sets the flag to 0.
RegSupp	Sets the flag to 1, if the register suppression timer is running, or else sets the flag to 0.
RegProbe	Sets the flag to 1, if the mcache entry is entering the register probing period, or else sets the flag to 0.
LSrc	Sets the flag to 1, if the source is in a directly-connected interface, or else sets the flag to 0.
LRcv	Sets the flag to 1, if the receiver is directly connected to the router, or else sets the flag to 0.
RP	Show the IP address of the RP.
num_oifs	Show the count of the outgoing interfaces.
immediate_oifs	Show the local immediate outgoing interface of the mcache entry.
inherited_oifs	Shows the PIM Sparse mode inherited outgoing interfaces.
blocked_oifs	Show the PIM Sparse mode blocked outgoing interfaces.
Flags	Show the flags associated with the forward entry.
fast ports ethe	Shows the forwarding port ID.
L3	Show whether the traffic is switched or routed out of the interface.
(HW)	Show whether the entry is software forwarded or hardware forwarded.
sm	Sets the flag to 1, if the Sparse Mode entry is enabled and 0, if the PIM dense mode entry is enabled.
ssm	Sets the flag to 1, if the SSM mode entry is enabled, or else sets the flag to 0.
hw	Sets the flag to 1, if the candidate for hardware forwarding is enabled, or else sets the flag to 0.
fast	Sets the flag to 1, if the resources are allocated for hardware forwarding, or else sets the flag to 0.
slow	Sets the flag to 1, if the entry is not a candidate for hardware forwarding or the resource allocation is failed, or else sets the flag to 0.
prun	Sets the flag to 1, if PIM is enabled, or else sets the flag to 0.

TABLE 19 Output parameters of the `show ip pim mcache mvid 1` command (continued)

Field	Description
tag	Sets the flag to 1, if there is a need for allocating entries from the replication table, or else sets the flag to 0.
needRte	Sets the flag to 1, if there is no route to the source and RP is available, or else sets the flag to 0.
msdp_adv	Sets the flag to 1, if RP is responsible for the source and must be advertised to its peers.
AgeSlTmsk	Shows the slot number on which Management Processor (MP) expects ingress traffic.
FID	Shows the FID resource allocated for a particular entry.
MVID	Shows the MVID resource allocated for a particular entry.
RegPkt	Shows the number of PIM register packet received.
AvgRate	Shows the average data traffic rate for the mcache entry.
profile	Shows the profile ID associated with the stream.
Number of matching entries	Shows the number of mcache entry.

Displaying the PIM multicast cache for FID

To display the PIM multicast cache for a specified fid, enter the following command at any CLI level

```
device# show ip pim mcache fid 800e
Total entries in mcache: 17
 1 (49.1.1.100, 229.0.0.1) in v49 (tag e4/9), Uptime 00:02:17 Rate 0
   Sparse Mode, RPT=0 SPT=1 Reg=0 L2Reg=1 RegSupp=0 RegProbe=0 LSrc=1 LRcv=1
   Source is directly connected. RP 192.168.1.1
   num_oifs = 1
   immediate_oifs = 0, inherited_oifs = 1, blocked_oifs = 0
   52,4/12(00:02:17/0) Flags:00000004
   fast ports ethe 4/12
   L3 (HW) 1: e4/12(VL52)
   Flags (0x7046c8e1)
     sm=1 ssm=0 hw=1 fast=1 slow=0 leaf=0 prun=0 tag=1 needRte=0 msdp_adv=1
   AgeSlTmsk=00000008, FID: 0x800e MVID: 1, RegPkt 0, AvgRate 0 profile: none
 2 (49.1.1.108, 229.0.0.1) in v49 (tag e4/9), Uptime 00:03:32 Rate 11257
   Sparse Mode, RPT=0 SPT=1 Reg=0 L2Reg=1 RegSupp=0 RegProbe=0 LSrc=1 LRcv=1
   Source is directly connected. RP 192.168.1.1
   num_oifs = 1
   immediate_oifs = 0, inherited_oifs = 1, blocked_oifs = 0
   52,4/12(00:03:32/0) Flags:00000004
   fast ports ethe 4/12
   L3 (HW) 1: e4/12(VL52)
   Flags (0x7046c8e1)
     sm=1 ssm=0 hw=1 fast=1 slow=0 leaf=0 prun=0 tag=1 needRte=0 msdp_adv=1
   AgeSlTmsk=00000008, FID: 0x800e MVID: 1, RegPkt 0, AvgRate 11257 profile: none
 3 (49.1.1.108, 229.0.0.2) in v49 (tag e4/9), Uptime 00:03:32 Rate 11257
   Sparse Mode, RPT=0 SPT=1 Reg=0 L2Reg=1 RegSupp=0 RegProbe=0 LSrc=1 LRcv=1
   Source is directly connected. RP 192.168.1.1
   num_oifs = 1
   immediate_oifs = 0, inherited_oifs = 1, blocked_oifs = 0
   52,4/12(00:03:32/0) Flags:00000004
   fast ports ethe 4/12
   L3 (HW) 1: e4/12(VL52)
   Flags (0x7046c8e1)
     sm=1 ssm=0 hw=1 fast=1 slow=0 leaf=0 prun=0 tag=1 needRte=0 msdp_adv=1
   AgeSlTmsk=00000008, FID: 0x800e MVID: 1, RegPkt 0, AvgRate 11257 profile: none
Number of matching entries: 3
```

Syntax: `show ip pim mcache fid fid-id`

The `fid-id` variable allows you to display an entry that matches a specified fid.

Table 20 describes the output parameters of the **show ip pim mcache fid 8** command.

TABLE 20 Output parameters of the **show ip pim mcache fid 8** command

Field	Description
Total entries in mcache	Shows the total number of PIM mcache entries.
v	Shows the interface name.
tag e	Shows the VLAN tagged ethernet interface.
Uptime	Shows the entry uptime.
Rate	Shows the total packet count.
Sparse mode	Shows that the PIM sparse mode is enabled.
RPT	Sets the flag to 1, if the entry is on the RP tree, else sets the flag to 0.
SPT	Sets the flag to 1, if the entry is on the source tree, or else sets the flag to 0.
Reg	Sets the flag to 1, if the data registration is in progress, or else sets the flag to 0.
L2Reg	Sets the flag to 1, if the source is directly connected to the router, or else sets the flag to 0.
RegSupp	Sets the flag to 1, if the register suppression timer is running, or else sets the flag to 0.
RegProbe	Sets the flag to 1, if the mcache entry is entering the register probing period, or else sets the flag to 0.
LSrc	Sets the flag to 1, if the source is in a directly-connected interface, or else sets the flag to 0.
LRcv	Sets the flag to 1, if the receiver is directly connected to the router, or else sets the flag to 0.
RP	Show the IP address of the RP.
num_oifs	Show the count of the outgoing interfaces.
inherited_oifs	Show the PIM Sparse mode inherited outgoing interfaces.
blocked_oifs	Show the PIM Sparse mode blocked outgoing interfaces.
Flags	Show the flags associated with the forward entry.
fast ports ethe	Shows the forwarding port ID.
L3	Show whether the traffic is switched or routed out of the interface.
(HW)	Show whether the entry is software forwarded or hardware forwarded.
sm	Sets the flag to 1, if the Sparse Mode entry is enabled and 0, if the PIM dense mode entry is enabled.
ssm	Sets the flag to 1, if the SSM mode entry is enabled, or else sets the flag to 0.
hw	Sets the flag to 1, if the candidate for hardware forwarding is enabled, or else sets the flag to 0.
fast	Sets the flag to 1, if the resources are allocated for hardware forwarding, or else sets the flag to 0.
slow	Sets the flag to 1, if the entry is not a candidate for hardware forwarding or the resource allocation is failed, or else sets the flag to 0.
prun	Sets the flag to 1, if PIM is enabled, or else sets the flag to 0.
tag	Sets the flag to 1, if there is a need for allocating entries from the replication table, or else sets the flag to 0.
needRte	Sets the flag to 1, if there is no route to the source and RP is available, or else sets the flag to 0.

TABLE 20 Output parameters of the `show ip pim mcache fid 8` command (continued)

Field	Description
msdp_adv	Sets the flag to 1, if RP is responsible for the source and must be advertised to its peers.
AgeSlitMsk	Shows the slot number on which MP expects ingress traffic.
FID	Shows the FID resource allocated for a particular entry.
MVID	Shows the MVID resource allocated for a particular entry.
RegPkt	Shows the number of PIM register packet received.
AvgRate	Shows the average data traffic rate for the mcache entry.
profile	Shows the profile ID associated with the stream.
Number of matching entries	Shows the number of mcache entry.

Syntax: `show ip pim mcache fid fid-id`

The `fid-id` variable allows you to display an entry that matches a specified fid.

Clearing the PIM forwarding cache

You can clear the PIM forwarding cache using the following command.

```
device# clear ip pim cache
```

Syntax: `clear ip pim [vrf vrf-name] cache`

Use the `vrf` option to clear the PIM forwarding cache for a VRF instance specified by the `vrf-name` variable.

Displaying PIM traffic statistics

To display PIM traffic statistics, enter the following command at any CLI level.

```
device(config)# show ip pim traffic
Port  HELLO    JOIN      PRUNE     ASSERT    REGISTER  REGISTER  BOOTSTRAP  CAND.  RP
      HELLO    JOIN      PRUNE     ASSERT    GRAFT (DM) STOP (SM)  MSGS (SM) ADV. (SM)
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
      Rx      Rx      Rx      Rx      Rx      Rx      Rx      Rx      Rx
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
e1/2  113      15841    13608     0         0         0         56       0
v52   113      13585    13122     0         13276    0         56       0
v57   0        0        0         0         0         0         0        0
v59   223      81345    14268     0         77760    0         0        0

Port  HELLO    JOIN      PRUNE     ASSERT    REGISTER  REGISTER  BOOTSTRAP  CAND.  RP
      HELLO    JOIN      PRUNE     ASSERT    GRAFT (DM) STOP (SM)  MSGS (SM) ADV. (SM)
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
      Tx      Tx      Tx      Tx      Tx      Tx      Tx      Tx      Tx
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
e1/2  111      0        0         0         0         0         1        0
v52   112      0        0         0         0         13276    55       0
v57   111      0        0         0         0         0         0        0
v59   112      0        0         0         0         0         56       0
device(config)#
```

Syntax: `show ip pim [vrf vrf-name] traffic`

The `vrf` option allows you to display the PIM traffic statistics for the VRF instance identified by the `vrf-name` variable.

NOTE

If you have configured interfaces for standard PIM (dense mode) on the device, statistics for these interfaces are listed first by the display.

Table 21 describes the output from this command.

TABLE 21 Output from the `show ip pim vrf traffic` command

This field...	Displays...
Port	The port or virtual interface on which the PIM interface is configured.
Hello	The number of PIM Hello messages sent or received on the interface.
J or P	The number of Join or Prune messages sent or received on the interface. NOTE Unlike PIM dense, PIM Sparse uses the same messages for Joins and Prunes.
Register	The number of Register messages sent or received on the interface.
RegStop	The number of Register Stop messages sent or received on the interface.
Assert	The number of Assert messages sent or received on the interface.
Total Recv or Xmit	The total number of IGMP messages sent and received by the device.
Total Discard or chksum	The total number of IGMP messages discarded, including a separate counter for those that failed the checksum comparison.

Clearing the PIM message counters

You can clear the PIM message counters using the following command.

```
device# clear ip pim traffic
```

Syntax: `clear ip pim [vrf vrf-name] traffic`

Use the **vrf** option to clear the PIM message counters for a VRF instance specified by the **vrf-name** variable.

Displaying PIM RPF

The `show ip pim rpf` command displays what PIM sees as the reverse path to the source as shown in the following. While there may be multiple routes back to the source, the one displayed by this command is the one that PIM thinks is best.

```
device# show ip pim vrf eng rpf 130.50.11.10
Source 130.50.11.10 directly connected on e4/1
```

Syntax: `show ip pim [vrf vrf-name] rpf ip-address`

The **ip-address** variable specifies the source address for RPF check.

The **vrf** option to display what PIM sees as the reverse path to the source for a VRF instance specified by the **vrf-name** variable.

Displaying PIM counters

You can display the number of default-vlan-id changes that have occurred since the applicable VRF was created, and how many times a tagged port was placed in a VLAN since the applicable VRF was created as shown.

```
device(config)# show ip pim vrf eng counter
Event Callback:
  DFTVlanChange      :           0           VlanPort      :           0
```

```

LP to MP IPCs:
SM_REGISTER      :      94924          MCAST_CREATE    :           0
S_G_AGEOUT      :           0          WRONG_IF        :      378153
ABOVE_THRESHOLD :       255          MCAST_FIRST_DATA :       765
MP to LP IPCs:
INIT            :       2317          INSERT_VPORT    :      4465
DELETE_VPORT   :       255          DELETE_VIF      :       255
MOVE_VPORT     :           0          DEL_ENTRY       :       510
INSERT_SOURCE  :           0          DELETE_SOURCE   :           0
RESET_SRC_LIST :           0          MOVE_TNNL_PORT  :           0
AGE_RESET      :      94924          FLAG_CHANGE     :       741
OIF_FLAG_CHANGE :       741
Error Counters:
RPSET_MAXED    :           0
device(config)#

```

Syntax: `show ip pim [vrf vrf-name] counter`

Table 22 describes the output from this command.

TABLE 22 Output from the `show ip pim vrf counter` command

This field..	Displays..
DFTVlanChange	The number of default-vlan-id changes that have occurred since the applicable VRF was created.
VlanPort	The number of times that a tagged port was placed in a VLAN since the applicable VRF was created.

NOTE

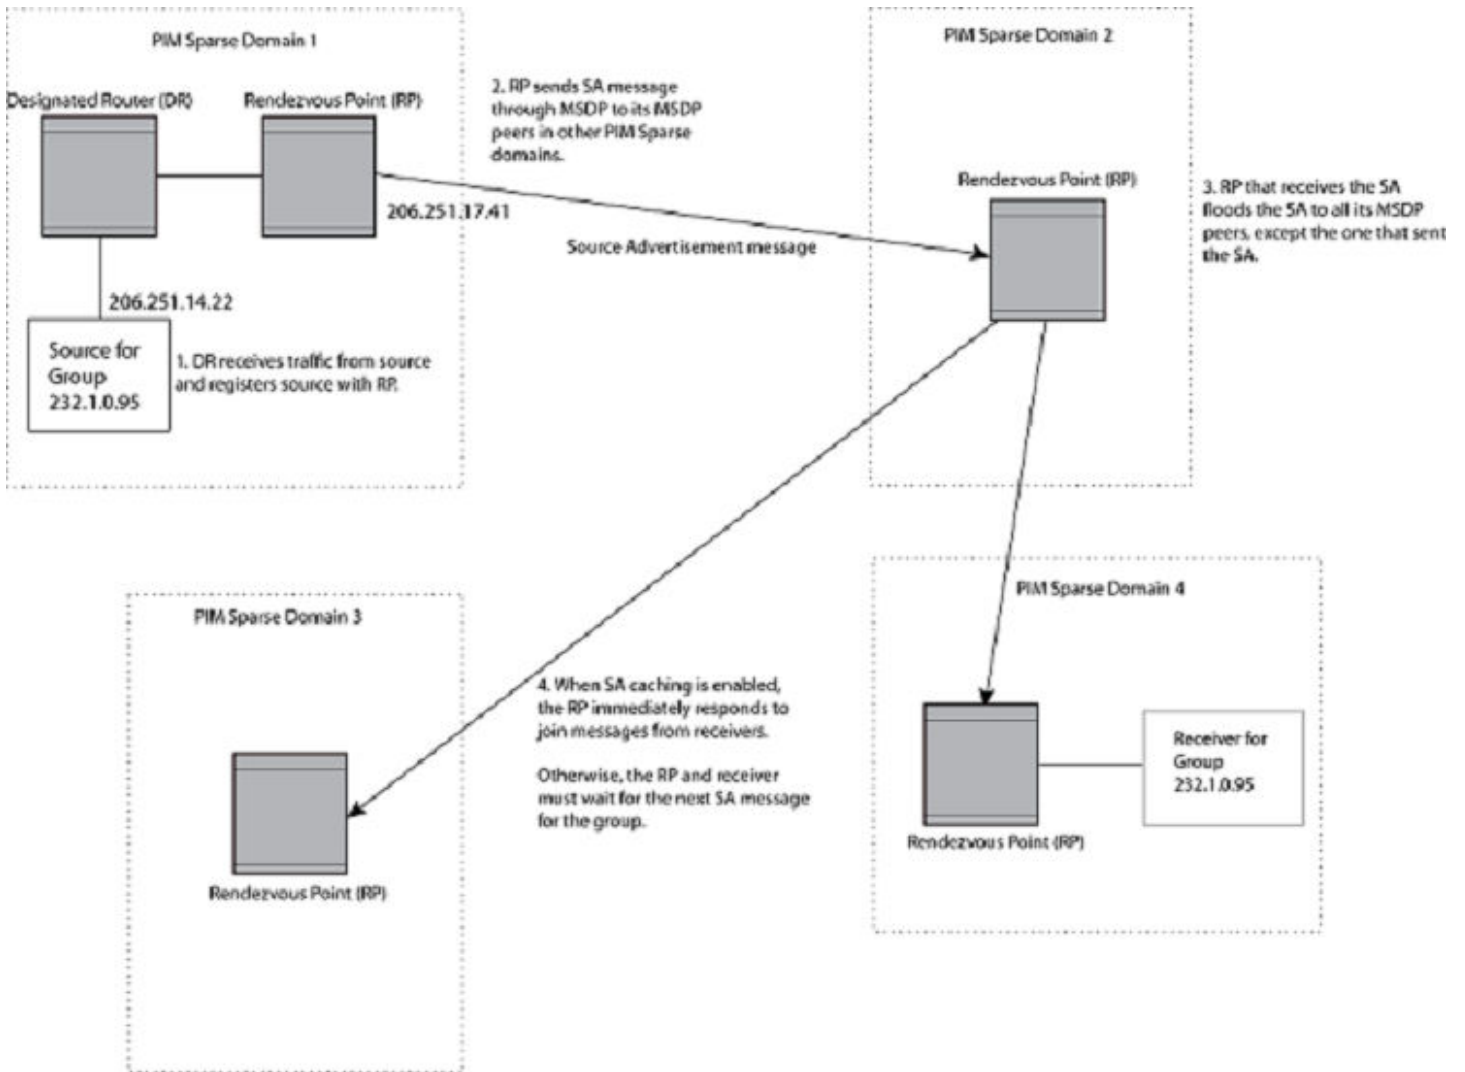
Since VLANs are not VRF-aware, any changes to default-vlan or tagged port moves is counted by all VRFs in existence at the time, including the default VRF.

Configuring Multicast Source Discovery Protocol (MSDP)

The Multicast Source Discovery Protocol (MSDP) is used by Protocol Independent Multicast (PIM) Sparse devices to exchange source information across PIM Sparse domains. Devices running MSDP can discover PIM Sparse sources in other PIM Sparse domains.

Figure 14 shows an example of some PIM Sparse domains. For simplicity, this example shows one Designated Router (DR), one group source, and one receiver for the group. Only one PIM Sparse device within each domain needs to run MSDP.

FIGURE 14 PIM Sparse domains joined by MSDP devices



In this example, the source for PIM Sparse multicast group 232.0.1.95 is in PIM Sparse domain 1. The source sends a packet for the group to its directly attached DR. The DR sends a Group Advertisement message for the group to the RP for the domain. The RP is configured for MSDP, which enables the RP to exchange source information with other PIM Sparse domains by communicating with RPs in other domains that are running MSDP.

The RP sends the source information to each peer through a Source Active message. The message contains the IP address of the source, the group address to which the source is sending, and the IP address of the RP.

In this example, the Source Active message contains the following information:

- Source address: 206.251.14.22
- Group address: 232.1.0.95
- RP address: 206.251.17.41

Figure 14 shows only one peer for the MSDP device (which is also the RP here) in domain 1, so the Source Active message goes to only that peer. When an MSDP device has multiple peers, it sends a Source Active message to each of those peers. Each peer sends the

Source Advertisement to other MSDP peers. The RP that receives the Source Active message also sends a Join message to the source if the RP that received the message has receivers for the group and source.

Peer Reverse Path Forwarding (RPF) flooding

When the MSDP device (also the RP) in domain 2 receives the Source Active message from the peer in domain 1, the MSDP device in domain 2 forwards the message to all other peers. This propagation process is sometimes called "peer Reverse Path Forwarding (RPF) flooding". In [Configuring Multicast Source Discovery Protocol \(MSDP\)](#) on page 107, the MSDP device floods the Source Active message it receives from the peer in domain 1 to peers in domains 3 and 4.

The MSDP device in domain 2 does not forward the Source Active back to the peer in domain 1, because that is the peer from which the device received the message. An MSDP device never sends a Source Active message back to the peer that sent it. The peer that sent the message is sometimes called the "RPF peer". The MSDP device uses the unicast routing table for its Exterior Gateway Protocol (EGP) to identify the RPF peer by looking for the route entry that is the next hop toward the source. Often, the EGP protocol is Border Gateway Protocol (BGP) version 4.

NOTE

MSDP depends on BGP and MBGP for inter-domain operations.

The MSDP routers in domains 3 and 4 also forward the Source Active message to all peers except the ones that sent them the message. [Configuring Multicast Source Discovery Protocol \(MSDP\)](#) on page 107 does not show additional peers.

Source Active caching

When an MSDP device that is also an RP and source receives a Source Active message, the RP and source checks the PIM Sparse multicast group table for receivers for the group. If the DR has a receiver for the group being advertised in the Source Active message, the RP sends a Join message towards that source.

In [Configuring Multicast Source Discovery Protocol \(MSDP\)](#) on page 107, if the MSDP device and RP in domain 4 has a table entry for the receiver, the RP sends a Join message on behalf of the receiver back through the RPF tree to the source, in this case the source in domain 1.

Source Active caching is enabled in MSDP on Extreme devices. The RP caches the Source Active messages it receives even if the RP does not have a receiver for the group. Once a receiver arrives, the RP can then send a Join to the cached source immediately.

The size of the cache used to store MSDP Source Active messages is 32K.

Configuring MSDP

To configure MSDP, perform the following tasks:

- Enable MSDP.
- Configure the MSDP peers.

NOTE

The PIM Sparse Rendezvous Point (RP) is also an MSDP peer.

NOTE

Devices that run MSDP usually also run BGP. The source address used by the MSDP device is normally configured to be the same source address used by BGP.

Enabling MSDP

To enable MSDP, enter the following command.

```
device(config)# router msdp
```

Syntax: [no] router msdp

Enabling MSDP for a specified VRF

The **vrf** parameter allows you to configure MSDP on the virtual routing instance (VRF) specified by the **vrf-name** variable. All MSDP parameters available for the default router instance are configurable for a VRF-based MSDP instance.

To enable MSDP for the VRF named "blue", enter the following commands.

```
device(config)# router msdp vrf blue
device(config-msdp-router-vrf-blue)
```

Syntax: [no] router msdp [vrf *vrf-name*]

The **vrf** parameter allows you to configure MSDP on the virtual routing instance (VRF) specified by the **vrf-name** variable.

Entering a **no router msdp vrf** command removes the MSDP configuration from the specified VRF only.

Configuring MSDP peers

To configure an MSDP peer, enter a command such as the following at the MSDP configuration level.

```
device(config-msdp-router)# msdp-peer 205.216.162.1
```

To configure an MSDP peer on a VRF, enter the following commands at the MSDP VRF configuration level.

```
device(config)# router msdp vrf blue
device(config-msdp-router-vrf-blue)# msdp-peer 205.216.162.1
```

Syntax: [no] msdp-peer *ip-addr* [connect-source loopback *num*]

The **ip-addr** parameter specifies the IP address of the neighbor.

The **connect-source loopbacknum** parameter specifies the loopback interface you want to use as the source for sessions with the neighbor and must be reachable within the VRF.

NOTE

It is strongly recommended that you use the connect-source loopback **num** parameter when issuing the **msdp-peer** command. If you do not use this parameter, the device uses the IP address of the outgoing interface. You should also make sure the IP address of the connect-source loopback is the source IP address used by the PIM-RP, and the BGP device.

The commands in the following example add an MSDP neighbor and specify a loopback interface as the source interface for sessions with the neighbor. By default, the device uses the subnet address configured on the physical interface where you configure the neighbor as the source address for sessions with the neighbor.

```
device(config)# interface loopback 1
device(config-lbif-1)# ip address 9.9.9.9/32
device(config)# router msdp
device(config-msdp-router)# msdp-peer 2.2.2.99 connect-source loopback 1
```

Disabling an MSDP peer

To disable an MSDP peer, enter the following command at the configure MSDP router level.

```
device(config-msdp-router)# msdp-peer 205.216.162.1 shutdown
```

To disable the MSDP VRF peer named "blue", enter the following commands.

```
device(config)# router msdp vrf blue
device(config-msdp-router-vrf-blue)# no msdp-peer 205.216.162.1
```

Syntax: `[no] msdp-peer ip-addr shutdown`

The **ip-addr** parameter specifies the IP address of the MSDP peer that you want to disable.

Designating the interface IP address as the RP IP address

When an RP receives a Source Active message, it checks its PIM Sparse multicast group table for receivers for the group. If a receiver exists the RP sends a Join to the source.

By default, the IP address included in the RP address field of the SA message is the IP address of the originating RP. An SA message can use the IP address of any interface on the originating RP. (The interface is usually a loopback interface.)

To designate an interface IP address to be the IP address of the RP, enter commands such as the following.

```
device(config)#
interface loopback 2
device(config-lbif-2)# ip address 2.2.1.99/32
device(config)# router msdp
device(config-msdp-router)# originator-id loopback 2
device(config-msdp-router)# exit
```

To specify VRF information, enter the following commands at the MSDP VRF configuration level.

```
device(config)#
interface loopback 2
device(config-lbif-2)# ip address 2.2.1.99/32
device(config)# router msdp vrf blue
device(config-msdp-router-vrf blue)# originator-id loopback 2
device(config-msdp-router-vrf blue)# exit
```

Syntax: `[no] originator-id type number`

The **originator-id** parameter instructs MSDP to use the specified interface IP address as the IP address of the RP in an SA message. This address must be the address of the interface used to connect the RP to the source. The default address used is the RP IP address.

The **type** parameter indicates the type of interface used by the RP. Ethernet, loopback and virtual routing interfaces (ve) can be used.

The **number** parameter specifies the interface number (for example: loopback number, port number or virtual routing interface number.)

Filtering MSDP source-group pairs

You can filter individual source-group pairs in MSDP Source-Active messages:

- **sa-filter in** - Filters source-group pairs received in Source-Active messages from an MSDP neighbor.
- **sa-filter originate** - Filters self-originated source-group pairs in outbound Source-Active messages sent to an MSDP neighbor
- **sa-filter out** - Filters self-originated and forwarded source-group pairs in outbound Source-Active messages sent to an MSDP neighbor

Filtering incoming and outgoing Source-Active messages

The following example configures filters for incoming Source-Active messages from three MSDP neighbors:

- For peer 2.2.2.99, all source-group pairs in Source-Active messages from the neighbor are filtered (dropped).
- For peer 2.2.2.97, all source-group pairs except those with source address matching 10.x.x.x and group address of 235.10.10.1 are permitted.
- For peer 2.2.2.96, all source-group pairs except those associated with RP 2.2.42.3 are permitted.

To configure filters for incoming Source-Active messages, enter commands at the MSDP VRF configuration level.

To configure filters for outbound Source-Active messages, enter the optional out keyword.

Example

The following commands configure extended ACLs. The ACLs will be used in route maps, which will be used by the Source-Active filters.

```
device(config)# access-list 123 permit ip 10.0.0.0 0.255.255.255 host 235.10.10.1
device(config)# access-list 124 permit ip host 2.2.42.3 any
device(config)# access-list 125 permit ip any any
```

The following commands configure the route maps.

```
device(config)# route-map msdp_map deny 1
device(config-routemap msdp_map)# match ip address 123
device(config-routemap msdp_map)# exit
device(config)# route-map msdp_map permit 2
device(config-routemap msdp_map)# match ip address 125
device(config-routemap msdp_map)# exit
device(config)# route-map msdp2_map permit 1
device(config-routemap msdp2_map)# match ip address 125
device(config-routemap msdp2_map)# exit
device(config)# route-map msdp2_rp_map deny 1
device(config-routemap msdp2_rp_map)# match ip route-source 124
device(config-routemap msdp2_rp_map)# exit
device(config)# route-map msdp2_rp_map permit 2
device(config-routemap msdp2_rp_map)# match ip route-source 125
device(config-routemap msdp2_rp_map)# exit
```

To specify VRF information, enter the following commands at the MSDP VRF configuration level.

```
device(config)# router msdp vrf blue
device(config-msdp-router-vrf blue)# sa-filter in 2.2.2.99
device(config-msdp-router-vrf blue)# sa-filter in 2.2.2.97 route-map msdp_map
device(config-msdp-router-vrf blue)# sa-filter in 2.2.2.96 route-map msdp2_map rp-route-map msdp2_rp_map
```

The **sa-filter** commands configure the following filters:

- **sa-filter in 2.2.2.99** - This command drops all source-group pairs received from neighbor 2.2.2.99.

NOTE

The default action is to deny all source-group pairs from the specified neighbor. If you want to permit some pairs, use route maps.

- **sa-filter in 2.2.2.97 route-map msdp_map** - This command drops source-group pairs received from neighbor 2.2.2.97 if the pairs have source addresses matching 10.x.x.x and group address 235.10.10.1.
- **sa-filter in 2.2.2.96 route-map msdp2_map rp-route-map msdp2_rp_map** - This command accepts all source-group pairs except those associated with RP 2.2.42.3.

Syntax: [no] sa-filter in | originate | out ip-addr [route-map map-tag] [rp-route-map rp-map-tag]

Selecting the in option applies the filter to incoming Source-Active messages.

Selecting the **originate** option applies the filter to self-originated outbound Source-Active messages.

Selecting the **out** option applies the filter to self-originated and forwarded outbound Source-Active messages.

The **ip-addr** parameter specifies the IP address of the MSDP neighbor. The filters apply to Source-Active messages received from or sent to this neighbor.

The **route-mapmap-tag** parameter specifies a route map. The device applies the filter to source-group pairs that match the route map. Use the **match ip addressacl-id** command in the route map to specify an extended ACL that contains the source addresses.

The **rp-route-maprp-map-tag** parameter specifies a route map to use for filtering based on Rendezvous Point (RP) address. Use this parameter if you want to filter Source-Active messages based on their originating RP. Use the **match ip route-sourceacl-id** command in the route map to specify an extended ACL that contains the RP address.

NOTE

The default filter action is deny. If you want to permit some source-group pairs, use a route map.

Filtering advertised Source-Active messages

The following example configures the device to advertise all source-group pairs except the ones that have source address 10.x.x.x.

The following commands configure extended ACLs to be used in the route map definition.

```
device(config)# access-list 123 permit ip 10.0.0.0 0.255.255.255 any
device(config)# access-list 125 permit ip any any
```

The following commands use the above ACLs to configure a route map which denies source-group with source address 10.x.x.x and any group address, while permitting everything else.

```
device(config)# route-map msdp_map deny 1
device(config-routemap msdp_map)# match ip address 123
device(config-routemap msdp_map)# exit
device(config)# route-map msdp_map permit 2
device(config-routemap msdp_map)# match ip address 125
device(config-routemap msdp_map)# exit
```

The following commands configure the Source-Active filter.

```
device(config)# router msdp
device(config-msdp-router)# sa-filter originate route-map msdp_map
```

To specify VRF information, enter the following commands at the MSDP VRF configuration level.

```
device(config)# router msdp vrf blue
device(config-msdp-router-vrf blue)# sa-filter originate route-map msdp_map
```

Syntax: [no] **sa-filter originate** [**route-map map-tag**]

The **route-mapmap-tag** parameter specifies a route map. The router applies the filter to source-group pairs that match the route map. Use the **match ip addressacl-id** command in the route map to specify an extended ACL that contains the source and group addresses.

NOTE

The default filter action is deny. If you want to permit some source-group pairs, use a route map. A permit action in the route map allows the device to advertise the matching source-group pairs. A deny action in the route map drops the source-group pairs from advertisements.

Displaying MSDP information

You can display the following MSDP information:

- **Summary information** - the IP addresses of the peers, the state of the device MSDP session with each peer, and statistics for keepalive, source active, and notification messages sent to and received from each of the peers
- **VRF Information** - Summary information for a specific VRF
- **Peer information** - the IP address of the peer, along with detailed MSDP and TCP statistics
- **Source Active cache entries** - the source active messages cached by the router.

Displaying summary information

To display summary MSDP information, enter the CLI command.

```
device(config)#show ip msdp vrf blue summary
MSDP Peer Status Summary
KA: Keepalive SA:Source-Active NOT: Notification
Peer Address Peer As State KA In Out SA In Out NOT In Out Age
40.40.40.1 1001 ESTABLISH 59 59 0 0 0 0 0 0 6
40.40.40.3 1001 ESTABLISH 59 59 0 0 0 0 0 0 47
47.1.1.2 N/A ESTABLISH 59 59 0 0 0 0 0 0 47
device(config)#
```

Syntax: show ip msdp summary

Table 23 describes the output from this command.

TABLE 23 MSDP summary information

This field...	Displays...
Peer address	The IP address of the peer interface with the device
State	The state of the MSDP device connection with the peer. The state can be one of the following: <ul style="list-style-type: none"> • CONNECTING - The session is in the active open state. • ESTABLISHED - The MSDP session is fully up. • INACTIVE - The session is idle. • LISTENING - The session is in the passive open state.
KA In	The number of MSDP keepalive messages the MSDP device has received from the peer
KA Out	The number of MSDP keepalive messages the MSDP device has sent to the peer
SA In	The number of source active messages the MSDP device has received from the peer
SA Out	The number of source active messages the MSDP device has sent to the peer
NOT In	The number of notification messages the MSDP router has received from the peer
NOT Out	The number of notification messages the MSDP router has sent to the peer

Displaying peer information

To display MSDP peer information, enter the following command.

```
device# show ip msdp vrf blue peer
Total number of MSDP Peers: 2
IP Address          State
1 206.251.17.30     ESTABLISHED
Keep Alive Time    Hold Time
60                 90
Message Sent      Message Received
Keep Alive        2                 3
Notifications     0                 0
Source-Active     0                 640
Last Connection Reset Reason:Reason Unknown
Notification Message Error Code Received:Unspecified
Notification Message Error SubCode Received:Not Applicable
Notification Message Error Code Transmitted:Unspecified
Notification Message Error SubCode Transmitted:Not Applicable
TCP Connection state: ESTABLISHED
Local host: 206.251.17.29, Local Port: 8270
Remote host: 206.251.17.30, Remote Port: 639
ISentSeq:         16927  SendNext:         685654  TotUnAck:         0
SendWnd:          16384  TotSent:          668727  ReTrans:          1
IRcvSeq:         45252428  RcvNext:         45252438  RcvWnd:           16384
TotalRcv:         10      RcvQueue:         0      SendQueue:        0
```

Syntax: show ip msdp peer

Table 24 describes the output from this command.

TABLE 24 MSDP peer information

This field...	Displays...
Total number of MSDP peers	The number of MSDP peers configured on the device
IP Address	The IP address of the peer's interface with the device
State	The state of the MSDP device connection with the peer. The state can be one of the following: <ul style="list-style-type: none"> CONNECTING - The session is in the active open state. ESTABLISHED - The MSDP session is fully up. INACTIVE - The session is idle. LISTENING - The session is in the passive open state.
Keep Alive Time	The keepalive time, which specifies how often this MSDP device sends keep alive messages to the neighbor. The keep alive time is 60 seconds and is not configurable.
Hold Time	The hold time, which specifies how many seconds the MSDP device will wait for a KEEPALIVE or UPDATE message from an MSDP neighbor before deciding that the neighbor is dead. The hold time is 75 seconds and is not configurable.
Keep Alive Message Sent	The number of keepalive messages the MSDP device has sent to the peer.
Keep Alive Message Received	The number of keepalive messages the MSDP device has received from the peer.
Notifications Sent	The number of notification messages the MSDP device has sent to the peer.
Notifications Received	The number of notification messages the MSDP device has received from the peer.
Source-Active Sent	The number of source active messages the MSDP device has sent to the peer.

TABLE 24 MSDP peer information (continued)

This field...	Displays...
Source-Active Received	The number of source active messages the MSDP device has received from the peer.
Last Connection Reset Reason	The reason the previous session with this neighbor ended.
Notification Message Error Code Received	<p>The MSDP device has received a notification message from the neighbor that contains an error code corresponding to one of the following errors. Some errors have subcodes that clarify the reason for the error. Where applicable, the subcode messages are listed underneath the error code messages:</p> <ul style="list-style-type: none"> • 1 - Message Header Error • 2 - SA-Request Error • 3 - SA-Message or SA-Response Error • 4 - Hold Timer Expired • 5 - Finite State Machine Error • 6 - Notification • 7 - Cease <p>For information about these errors, refer to section 17 in the Internet draft describing MSDP, "draft-ietf-msdp-spec".</p>
Notification Message Error SubCode Received	See above.
Notification Message Error Code Transmitted	The error message corresponding to the error code in the NOTIFICATION message this MSDP router sent to the neighbor. See the description for the Notification Message Error Code Received field for a list of possible codes.
Notification Message Error SubCode Transmitted	See above.
TCP Statistics	
TCP connection state	<p>The state of the connection with the neighbor. Can be one of the following:</p> <ul style="list-style-type: none"> • LISTEN - Waiting for a connection request. • SYN-SENT - Waiting for a matching connection request after having sent a connection request. • SYN-RECEIVED - Waiting for a confirming connection request acknowledgment after having both received and sent a connection request. • ESTABLISHED - Data can be sent and received over the connection. This is the normal operational state of the connection. • FIN-WAIT-1 - Waiting for a connection termination request from the remote TCP, or an acknowledgment of the connection termination request previously sent. • FIN-WAIT-2 - Waiting for a connection termination request from the remote TCP. • CLOSE-WAIT - Waiting for a connection termination request from the local user. • CLOSING - Waiting for a connection termination request acknowledgment from the remote TCP. • LAST-ACK - Waiting for an acknowledgment of the connection termination request previously sent to the remote TCP (includes an acknowledgment of the connection termination request). • TIME-WAIT - Waiting for enough time to pass to be sure the remote TCP received the acknowledgment of the connection termination request. • CLOSED - There is no connection state.

TABLE 24 MSDP peer information (continued)

This field...	Displays...
Local host	The IP address of the MSDP device interface with the peer.
Local port	The TCP port the MSDP router is using for the BGP4 TCP session with the neighbor.
Remote host	The IP address of the neighbor.
Remote port	The TCP port number of the peer end of the connection.
ISentSeq	The initial send sequence number for the session.
SendNext	The next sequence number to be sent.
TotUnAck	The number of sequence numbers sent by the MSDP device that have not been acknowledged by the neighbor.
SendWnd	The size of the send window.
TotSent	The number of sequence numbers sent to the neighbor.
ReTrans	The number of sequence numbers the MSDP device retransmitted because they were not acknowledged.
IRcvSeq	The initial receive sequence number for the session.
RcvNext	The next sequence number expected from the neighbor.
RcvWnd	The size of the receive window.
TotalRcv	The number of sequence numbers received from the neighbor.
RcvQue	The number of sequence numbers in the receive queue.
SendQue	The number of sequence numbers in the send queue.

Displaying Source Active cache information

To display the Source Actives in the MSDP cache, use the following command.

```
device# show ip msdp vrf blue sa-cache
Total of 10 SA cache entries
Index  RP address      (Source, Group)      Orig Peer      Age
1      2.2.2.2          (192.6.1.10, 227.1.1.1) 192.1.1.2     0
2      2.2.2.2          (192.6.1.10, 227.1.1.2) 192.1.1.2     0
3      2.2.2.2          (192.6.1.10, 227.1.1.3) 192.1.1.2     0
4      2.2.2.2          (192.6.1.10, 227.1.1.4) 192.1.1.2     0
5      2.2.2.2          (192.6.1.10, 227.1.1.5) 192.1.1.2     0
6      2.2.2.2          (192.6.1.10, 227.1.1.6) 192.1.1.2     0
7      2.2.2.2          (192.6.1.10, 227.1.1.7) 192.1.1.2     0
8      2.2.2.2          (192.6.1.10, 227.1.1.8) 192.1.1.2     0
9      2.2.2.2          (192.6.1.10, 227.1.1.9) 192.1.1.2     0
10     2.2.2.2          (192.6.1.10, 227.1.1.10) 192.1.1.2     0
```

Syntax: `show ip msdp sa-cache [source-address | group-address | peer-as as-number | counts | orig-rp rp-address | peer peer-address] [rejected | self-originated]`

The **source-address** parameter selects the source address of the SA entry.

The **group-address** parameter selects the group address of the SA entry.

The **peer-as** keyword specifies the BGP AS Number of the forwarding peer.

The **counts** keyword displays only the count of entries.

The **orig-rp** keyword specifies the originating RP address.

The **peer** keyword specifies the peer address.

The **rejected** keyword displays the rejected SAs.

The **self-originated** keyword displays the self-originated SAs.

Table 25 describes the output from this command.

TABLE 25 MSDP source active cache

This field...	Displays...
Total	The number of entries the cache currently contains.
Index	The cache entry number.
RP	The RP through which receivers can access the group traffic from the source
SourceAddr	The IP address of the multicast source.
GroupAddr	The IP multicast group to which the source is sending information.
Orig Peer	The peer from which this source-active entry was received.
Age	The number of seconds the entry has been in the cache

You can use the following command to filter the output to display only the entries matching a specific source.

```
device#show ip msdp sa-cache 1.1.1.1
```

You can use the following command to filter the output to display only the entries matching a specific group.

```
device#show ip msdp sa-cache 239.1.1.1
```

You can use the following command to filter the output to display only the SA cache entries that are received from peers in the BGP AS Number 100.

```
device#show ip msdp sa-cache 100
```

You can use the following command to filter the output to display only the SA cache entries that are originated by the RP 1.1.1.1.

```
device#show ip msdp sa-cache orig-rp 1.1.1.1
```

You can use the following command to filter the output to display only the SA cache entries that are received from the peer 1.1.1.1.

```
device#show ip msdp sa-cache peer 1.1.1.1
```

You can use the following command to display the rejected SAs. You can further narrow down by quoting the reason for rejection.

```
device#show ip msdp sa-cache rejected
```

You can use the following command to display the self-originated SAs.

```
device#show ip msdp sa-cache self-originated
```

Displaying MSDP RPF-Peer

To display MSDP peer information for the RP 1.1.1.1, enter the following command..

```
device# show ip msdp rpf-peer 1.1.1.1
MSDP Peer Status Summary
```

```
KA: Keepalive SA:Source-Active NOT: Notification
Peer Address Peer As State KA SA NOT Age
In Out In Out In Out
40.40.40.3 1001 ESTABLISH 62 62 0 0 0 0 7
device
```

Syntax: `show ip msdp rpf-peer ip-addr`

Displaying MSDP Peer

To display MSDP peer information, enter the following command.

```
device# show ip msdp peer 40.40.40.3
MSDP Peer Status Summary
KA: Keepalive SA:Source-Active NOT: Notification
Peer Address Peer As State KA SA NOT Age
In Out In Out In Out
40.40.40.3 1001 ESTABLISH 62 62 0 0 0 0 7
device#
```

Syntax: `show ip msdp peer peer-addr`

Displaying MSDP VRF RPF-Peer

To display MSDP peer information for a specific VRF, enter the following command.

```
device#sh ip msdp vrf Blue rpf-peer 40.40.40.2
MSDP Peer Status Summary
KA: Keepalive SA:Source-Active NOT: Notification
Peer Address Peer As State KA SA In Out In Out In Out
40.40.40.2 1001 ESTABLISH 5569 5568 0 0 0 0 57
```

Syntax: `show ip msdp vrf VRF-name rpf-peer ip-addr`

Clearing MSDP information

You can clear the following MSDP information:

- Peer information
- Source active cache
- MSDP statistics

Clearing peer information

To clear MSDP peer information, enter the following command at the Privileged EXEC level of the CLI.

```
device# clear ip msdp peer 205.216.162.1
```

Syntax: `clear ip msdp peer ip-addr`

The command in this example clears the MSDP peer connection with MSDP router 205.216.162.1. The CLI displays a message to indicate when the connection has been successfully closed. To clear all the peers, omit the `ip-addr` variable from the command.

Clearing peer information on a VRF

To clear the MSDP VRF peers, enter the following command at the MSDP VRF configuration level.

```
device#clear ip msdp vrf blue peer 207.207.162.5
```

Clearing the source active cache

To clear the source active cache, enter the following command at the Privileged EXEC level of the CLI.

```
device# clear ip msdp sa-cache
```

Syntax: `clear ip msdp sa-cache ip-addr`

The command in this example clears all the cache entries. Use the **ip-addr** variable to clear only the entries matching either a source or a group.

Clearing the source active cache for a VRF

To clear the MSDP VRF source active cache by entering the following command at the MSDP VRF configuration level.

```
device#clear ip msdp
vrf blue sa-cache
```

Syntax: `clear ip msdp [vrf vrf-name | ip-addr] sa-cache`

Clearing MSDP statistics

To clear MSDP statistics, enter the following command at the Privileged EXEC level of the CLI.

```
device# clear ip msdp statistics
```

Syntax: `clear ip msdp statistics ip-addr`

The command in this example clears statistics for all the peers. To clear statistics for only a specific peer, enter the IP address of the peer.

Clearing MSDP VRF statistics

To clear the MSDP VRF statistics by entering the following command.

```
device# clear ip msdp vrf blue statistics
```

Syntax: `clear ip msdp statistics [vrf vrf-name] [ip-addr]`

The command in this example clears statistics for all the peers. To clear statistics for only a specific peer, enter the IP address of the peer.

The command in this example clears all statistics for all the peers in the VRF "blue".

Configuring MSDP mesh groups


A PIM Sparse domain can have several RPs that are connected to each other to form an MSDP mesh group. To qualify as a mesh group, the RPs have to be fully meshed; that is, each RP must be connected to all peer RPs in a domain. (Refer to [Figure 15](#).)

A mesh group reduces the forwarding of SA messages within a domain. Instead of having every RP in a domain forward SA messages to all the RPs within that domain, only one RP forwards the SA message. Since an MSDP mesh group is fully meshed, peers do not forward SA messages received in a domain from one member to any member of the group. The RP that originated the SA or the first RP

in a domain that receives the SA message is the only one that forwards the message to the members of a mesh group. An RP can forward an SA message to any MSDP router as long as that peer is farther away from the originating RP than the current MSDP router.

Figure 15 shows an example of an MSDP mesh group. In a PIM-SM mesh group the RPs are configured to be peers of each other. They can also be peers of RPs in other domains.

FIGURE 15 Example of MSDP mesh group



PIM Sparse Domain 1 in Figure 15 contains a mesh group with four RPs. When the first RP, for example, RP 206.251.21.31 originates or receives an SA message from a peer in another domain, it sends the SA message to its peers within the mesh group. However, the

peers do not send the message back to the originator RP or to each other. The RPs then send the SA message farther away to their peers in other domains. The process continues until all RPs within the network receive the SA message.

Configuring MSDP mesh group

To configure an MSDP mesh group, enter commands such as the following on each device that will be included in the mesh group.

```
device(config)# router msdp
device(config-msdp-router)# msdp-peer 206.251.18.31 connect-source loopback 2
device(config-msdp-router)# msdp-peer 206.251.19.31 connect-source loopback 2
device(config-msdp-router)# msdp-peer 206.251.20.31 connect-source loopback 2
device(config-msdp-router)# mesh-group GroupA 206.251.18.31
device(config-msdp-router)# mesh-group GroupA 206.251.19.31
device(config-msdp-router)# mesh-group GroupA 206.251.20.31
device(config-msdp-router)# exit
```

Syntax: `[no] mesh-group group-name peer-address`

The sample configuration above reflects the configuration in [Configuring MSDP mesh groups](#) on page 120. On RP 206.251.21.31 you specify its peers within the same domain (206.251.18.31, 206.251.19.31, and 206.251.20.31).

You first configure the MSDP peers using the **msdp-peer** command to assign their IP addresses and the loopback interfaces.

Next, place the MSDP peers within a domain into a mesh group. Use the **mesh-group** command. There are no default mesh groups.

The **group-name** parameter identifies the mesh group. Enter up to 31 characters for group-name. You can have up to 4 mesh groups within a multicast network. Each mesh group can include up to 15 peers.

The **peer-address** parameter specifies the IP address of the MSDP peer that is being placed in the mesh group.

NOTE

On each of the device that will be part of the mesh group, there must be a mesh group definition for all the peers in the mesh-group.

A maximum of 15 MSDP peers can be configured per mesh group.

MSDP Anycast RP

MSDP Anycast RP is a method of providing intra-domain redundancy and load-balancing between multiple Rendezvous Points (RP) in a Protocol Independent Multicast Sparse mode (PIM-SM) network. It is accomplished by configuring all RPs within a domain with the same anycast RP address which is typically a loopback IP address. Multicast Source Discovery Protocol (MSDP) is used between all of the RPs in a mesh configuration to keep all RPs in sync regarding the active sources.

PIM-SM routers are configured to register (statically or dynamically) with the RP using the same anycast RP address. Since multiple RPs have the same anycast address, an Interior Gateway Protocol (IGP) such as OSPF routes the PIM-SM router to the RP with the best route. If the PIM-SM routers are distributed evenly throughout the domain, the loads on RPs within the domain will be distributed. If the RP with the best route goes out of service, the PIM-SM router's IGP changes the route to the closest operating RP that has the same anycast address.

This configuration works because MSDP is configured between all of the RPs in the domain. Consequently, all of the RPs share information about active sources.

This feature uses functionality that is already available on the Extreme device but re-purposes it to provide the benefits desired as described in RFC 3446.

Configuring MSDP Anycast RP

To configure MSDP Anycast RP, you must perform the following tasks:

- Configure a loopback interface with the anycast RP address on each of the RPs within the domain and enable PIM-SM on these interfaces.
- Ensure that the anycast RP address is leaked into the IGP domain. This is typically done by enabling the IGP on the loopback interface (in passive mode) or redistributing the connected loopback IP address into the IGP.

NOTE

The anycast RP address *must* not be the IGP router-id.

- Enable PIM-SM on all interfaces on which multicast routing is desired.
- Enable an IGP on each of the loopback interfaces and physical interfaces configured for PIM-SM.
- Configure loopback interfaces with unique IP addresses on each of the RPs for MSDP peering. This loopback interface is also used as the MSDP originator-id.
- The non-RP PIM-SM routers may be configured to use the anycast RP address statically or dynamically (by the PIMv2 bootstrap mechanism).

Example


The example shown in [Figure 16](#) is a simple MSDP Anycast-enabled network with two RPs and two PIM-SM routers. Loopback 1 in RP 1 and RP 2 have the same IP address. Loopback 2 in RP1 and Loopback 2 in RP2 have different IP addresses and are configured as MSDP peering IP addresses in a mesh configuration.

In the PIM configuration for PIM-SM routers PIMR1 and PIMR2 the RP address is configured to be the anycast RP address that was configured on the Loopback 1 interfaces on RP1 and RP2. OSPF is configured as the IGP for the network and all of the devices are in OSPF area 0.

Since PIMR1 has a lower cost path to RP1 and PIMR2 has a lower cost path to RP2 they will register with the respective RPs when both are up and running. This shares the load between the two RPs. If one of the RPs fails, the higher-cost path to the IP address of Loopback 1 on the RPs is used to route to the still-active RP.

The configuration examples demonstrate the commands required to enable this application.

FIGURE 16 Example of a MSDP Anycast RP network



RP 1 configuration

The following commands provide the configuration for the RP 1 router.

```
RP1(config)#router ospf
RP1(config-ospf-router)# area 0
RP1(config-ospf-router)# exit
RP1(config)# interface loopback 1
RP1(config-lbif-1)# ip ospf area 0
RP1(config-lbif-1)# ip ospf passive
RP1(config-lbif-1)# ip address 10.0.0.1/32
RP1(config-lbif-1)# ip pim-sparse
RP1(config-lbif-1)# exit
RP1(config)# interface loopback 2
RP1(config-lbif-2)# ip ospf area 0
RP1(config-lbif-2)# ip ospf passive
RP1(config-lbif-2)# ip address 10.1.1.1/32
RP1(config-lbif-2)# exit
RP1(config)# interface ethernet 5/1
RP1(config-if-e1000-5/1)# ip ospf area 0
RP1(config-if-e1000-5/1)# ip address 192.1.1.1/24
RP1(config-if-e1000-5/1)# ip pim-sparse
RP1(config)# interface ethernet 5/2
RP1(config-if-e1000-5/2)# ip ospf area 0
RP1(config-if-e1000-5/2)# ip ospf cost 5
RP1(config-if-e1000-5/2)# ip address 192.2.1.1/24
RP1(config-if-e1000-5/2)# ip pim-sparse
RP1(config)# interface ethernet 5/3
RP1(config-if-e1000-5/3)# ip ospf area 0
```

```

RP1(config-if-e1000-5/3)# ip ospf cost 10
RP1(config-if-e1000-5/3)# ip address 192.3.1.1/24
RP1(config-if-e1000-5/3)# ip pim-sparse
RP1(config-if-e1000-5/3)# exit
RP1(config)# router pim
RP1(config-pim-router)# rp-candidate loopback 1
RP1(config-pim-router)# exit
RP1(config)# router msdp
RP1(config-msdp-router)# msdp-peer 10.1.1.2 connect-source loopback 2
RP1(config-msdp-router)# originator-id loopback 2

```

RP 2 configuration

The following commands provide the configuration for the RP 2 router.

```

RP2(config)#router ospf
RP2(config-ospf-router)# area 0
RP2(config-ospf-router)# exit
RP2(config)# interface loopback 1
RP2(config-lbif-1)# ip ospf area 0
RP2(config-lbif-1)# ip ospf passive
RP2(config-lbif-1)# ip address 10.0.0.1/32
RP2(config-lbif-1)# ip pim-sparse
RP2(config-lbif-1)# exit
RP2(config)# interface loopback 2
RP2(config-lbif-2)# ip ospf area 0
RP2(config-lbif-2)# ip ospf passive
RP2(config-lbif-2)# ip address 10.1.1.2/32
RP2(config-lbif-2)# exit
RP2(config)# interface ethernet 5/1
RP2(config-if-e1000-5/1)# ip ospf area 0
RP2(config-if-e1000-5/1)# ip address 192.1.1.2/24
RP2(config-if-e1000-5/1)# ip pim-sparse
RP2(config)# interface ethernet 5/2
RP2(config-if-e1000-5/2)# ip ospf area 0
RP2(config-if-e1000-5/2)# ip ospf cost 5
RP2(config-if-e1000-5/2)# ip address 192.5.2.1/24
RP2(config-if-e1000-5/2)# ip pim-sparse
RP2(config)# interface ethernet 5/3
RP2(config-if-e1000-5/3)# ip ospf area 0
RP2(config-if-e1000-5/3)# ip ospf cost 10
RP2(config-if-e1000-5/3)# ip address 192.6.1.2/24
RP2(config-if-e1000-5/3)# ip pim-sparse
RP2(config-if-e1000-5/3)# exit
RP2(config)# router pim
RP2(config-pim-router)# rp-candidate loopback 1
RP2(config-pim-router)# exit
RP2(config)# router msdp
RP2(config-msdp-router)# msdp-peer 10.1.1.1 connect-source loopback 2
RP2(config-msdp-router)# originator-id loopback 2

```

PIMR1 configuration

The following commands provide the configuration for the PIMR1 router.

```

PIMR1(config)#router ospf
PIMR1(config-ospf-router)# area 0
PIMR1(config-ospf-router)# exit
PIMR1(config)# interface ethernet 6/2
PIMR1(config-if-e1000-6/2)# ip ospf area 0
PIMR1(config-if-e1000-6/2)# ip ospf cost 5
PIMR1(config-if-e1000-6/2)# ip address 192.2.1.2/24
PIMR1(config-if-e1000-6/2)# ip pim-sparse
PIMR1(config)# interface ethernet 6/3
PIMR1(config-if-e1000-6/3)# ip ospf area 0
PIMR1(config-if-e1000-6/3)# ip ospf cost 10
PIMR1(config-if-e1000-6/3)# ip address 192.6.1.1/24
PIMR1(config-if-e1000-6/3)# ip pim-sparse

```

```
PIMR1(config-if-e1000-6/3)# exit
PIMR1(config)# router pim
PIMR1(config-pim-router)# rp-address 10.0.0.1
PIMR1(config-pim-router)# exit
```

PIMR2 configuration

The following commands provide the configuration for the PIMR2 router.

```
PIMR2(config)#router ospf
PIMR2(config-ospf-router)# area 0
PIMR2(config-ospf-router)# exit
PIMR2(config)# interface ethernet 1/2
PIMR2(config-if-e1000-1/2)# ip ospf area 0
PIMR2(config-if-e1000-1/2)# ip ospf cost 5
PIMR2(config-if-e1000-1/2)# ip address 192.5.2.2/24
PIMR2(config-if-e1000-1/2)# ip pim-sparse
PIMR2(config)# interface ethernet 1/3
PIMR2(config-if-e1000-1/3)# ip ospf area 0
PIMR2(config-if-e1000-1/3)# ip ospf cost 10
PIMR2(config-if-e1000-1/3)# ip address 192.3.1.2/24
PIMR2(config-if-e1000-1/3)# ip pim-sparse
PIMR2(config-if-e1000-1/3)# exit
PIMR2(config)# router pim
PIMR2(config-pim-router)# rp-address 10.0.0.1
PIMR2(config-pim-router)# exit
```

PIM Anycast RP

PIM Anycast RP is a method of providing load balancing and fast convergence to PIM RPs in an IPv4 multicast domain. The RP address of the Anycast RP is a shared address used among multiple PIM routers, known as PIM RP. The PIM RP routers create an Anycast RP set. Each router in the Anycast RP set is configured using two IP addresses; a shared RP address in their loopback address and a separate, unique ip address. The loopback address must be reachable by all PIM routers in the multicast domain. The separate, unique ip address is configured to establish static peering with other PIM routers and communication with the peers.

When the source is activated in a PIM Anycast RP domain, the PIM First Hop (FH) will register the source to the closet PIM RP. The PIM RP follows the same MSDP Anycast RP operation by decapsulating the packet and creating the (s,g) state. If there are external peers in the Anycast RP set, the router will re-encapsulate the packet with the local peering address as the source address of the encapsulation. The router will unicast the packet to all Anycast RP peers. The re-encapsulation of the data register packet to Anycast RP peers ensures source state distribution to all RPs in a multicast domain.

Configuring PIM Anycast RP

A new PIM CLI is introduced for PIM Anycast RP under the router pim sub mode. The PIM CLI specifies mapping of the RP and the Anycast RP peers.

To configure PIM Anycast RP, enter the following command.

```
device(config)#router pim
device(config-pim-router)#rp-address 100.1.1.1
device(config-pim-router)#anycast-rp 100.1.1.1 my-anycast-rp-set-acl
```

Syntax: [no] **anycast-rp** *rp-address* **anycast-rp-set-acl**

The **rp address** parameter specifies a shared RP address used among multiple PIM routers.

The **anycast-rp-set-acl** parameter specifies a host based simple acl used to specifies the address of the Anycast RP set, including a local address.

The following example is a configuration of PIM Anycast RP 100.1.1.1. The example avoids using loopback 1 interface when configuring PIM Anycast RP because the loopback 1 address could be used as a router-id. A PIM First Hop router will register the source with the closest RP. The first RP that receives the register will re-encapsulate the register to all other Anycast RP peers. Please refer to [Figure 17](#) as described in the configuration of PIM Anycast RP 100.1.1.1.

```
device(config)#interface loopback 2
device(config-lbif-2)#ip address 100.1.1.1/24
device(config-lbif-2)#ip pim-sparse
device(config-lbif-2)#interface loopback 3
device(config-lbif-3)#ip address 1.1.1.1/24
device(config-lbif-3)#ip pim-sparse
device(config-lbif-3)#router pim
device(config-pim-router)#rp-address 100.1.1.1
device(config-pim-router)#anycast-rp 100.1.1.1 my-anycast-rp-set
device(config-pim-router)#ip access-list standard my-anycast-rp-set
device(config-std-nacl)#permit host 1.1.1.1
device(config-std-nacl)#permit host 2.2.2.2
device(config-std-nacl)#permit host 3.3.3.3
```


The RP shared address 100.1.1.1 is used in the PIM domain. IP addresses 1.1.1.1, 2.2.2.2, and 3.3.3.3 are listed in the ACL that forms the self-inclusive Anycast RP set. Multiple anycast-rp instances can be configured on a system; each peer with the same or different Anycast RP set.

NOTE

The PIM software supports up to eight PIM Anycast-RP routers. All deny statements in the anycast_rp_set acl and additional routers more than eight listed in an access list are ignored.

The example shown in [Figure 17](#) is a PIM Anycast-enabled network with 3 RPs, 1 PIM-FH router connecting to its active source and local receiver. Loopback 2 in RP1, RP2, and RP3 have the same IP addresses 100.1.1.1. Loopback 3 in RP1, RP2, and RP3 each have separate IP addresses configured to communicate with their peers in the Anycast RP set.

FIGURE 17 Example of a PIM Anycast RP network



Displaying information for a PIM Anycast RP interface

To display information for a PIM Anycast RP interface, enter the following command.

```
device(config)#show ip pim anycast-rp
Number of Anycast RP: 1
Anycast RP: 100.1.1.1
ACL ID: 200
ACL Name: my-anycast-rp-set
ACL Filter: SET
Peer List:
  1.1.1.1
```

2.2.2.2
3.3.3.3

Syntax: show ip pim anycast-rp

The following table describes the parameters of the **show ip pim anycast-rp** command:

TABLE 26 Display of show ip pim-anycast-rp

This field...	Displays...
Number of Anycast RP:	The Number of Anycast RP specifies the number of Anycast RP sets in the multicast domain.
Anycast RP:	The Anycast RP address specifies a shared RP address used among multiple PIM routers.
ACL ID:	The ACL ID specifies the ACL ID assigned.
ACL Name	The ACL Name specifies the name of the Anycast RP set.
ACL Filter	The ACL Filter specifies the ACL filter state SET or UNSET.
Peer List	The Peer List specifies host addresses that are permitted in the Anycast RP set.

NOTE

MSDP and Anycast RP do not interoperate. If transitioning from MSDP to Anycast RP or vice versa, all RPs in the network must be configured for the same method of RP peering; either Anycast RP or MSDP.

PIM over MCT intermediate router functionality

MCT peers support intermediate router functionality by accepting PIM neighbors on specific interfaces, thus routing multicast traffic as fully functional PIM devices acting as upstream and downstream routers.

NOTE

For additional information about PIM over MCT, refer to the *Brocade NetIron Switching Configuration Guide*.


MCT peers support Cluster Client Edge Port (CCEP) and Cluster Edge Port (CEP) interfaces.

PIM states between the MCT peers are synchronized by sending the control packets via MAC Database Update Protocol (MDUP). This is required due to the nature of the MCT LAG. Packets from the MCT client on the CCEP ports are received by only one of the MCT peers. Hence the control packets that are received natively on the CCEP ports are sent via MDUP to synchronize the states. The Join or Prune and Asserts are synchronized to maintain the Outgoing Interface (OIF) state for the CCEP ports on both peers. For OIFs created by PIM joins, only one of the MCT peers will forward the traffic and the other peer will keep the OIF in the MCT blocked state.

These are the general rules followed for the control packet handling algorithm.

- Control packets originated from MCT peers will be flooded on MCT VLAN. Exceptions are Assert packets.
- Control packets received on CCEP ports will be synchronized via MDUP being sent on CEP and other local CCEP ports. Exceptions are BSM packets, join and prune.
- Control packets received on CEP ports will be flooded on MCT vlan. Exceptions are Join and Prune.
- Control packets received on ICL will be flooded in controlled manner on MCT VLAN based on remote CCEP status i.e. based on whether they are up or down. Exceptions are join and prune.
- Control packets received on MDUP will be sent on CEP ports and other local CCEP ports based on remote CCEP being up or down.

FIGURE 18 PIM over MCT reference topology




MCT peer as intermediate Upstream router

P1 and P2 are the MCT peers and are acting as upstream routers for R1. R1 is the last-hop router (LHR).

P1, P2, and R1 are configured with PIM on the MCT virtual Ethernet (VE) interface. RP and source is in the core and the connectivity to the core is via an uplink.

FIGURE 19 MCT peer as immediate Upstream router



Hello exchange and neighbor state:

- In MCT topology, the CCEP links going out of P1 and P2 to R1 are treated as a single LAG at R1. Hence when R1 sends multicast packets (irrespective of control packets or data packets) they will reach only one of the peers. These control packets (Hellos, join or prune etc.) received by one peer will be sent to the other peer via the MDUP channel.
- Hellos sent by R1 could reach either P1 or P2 due to the above nature of MCT LAG.
- These Hellos that are reaching P2 are sent to P1 via MDUP channel. Thus P1 learns about R1 and treats this Hello as if it is received on its CCEP interface. Thus both P1 and P2 learn about the PIM neighbors across the CCEP links and create neighbor state for R1.
- Hellos originated from P1 and P2 are flooded on the MCT VLAN i.e. on ICL, CEP, local CCEP ports. This enables R1 to learn that both the MCT peers are PIM neighbors and also enables P1 and P2 to learn about each other as PIM neighbors on an ICL link and create neighbor state, for each other.

Join or prune exchange and mcache state:

- As receivers are connected to R1, R1 creates *,G state and sends a join state towards RP and sends it on the MCT LAG. This join, like any other packet, is received by only one of the MCT peers.
- Suppose P2 receives the *,G join natively. This join will be processed or consumed and will also be sent to P1 via MDUP.
- P1 processes this join received via MDUP as if it is received on CCEP.
- Both P1 and P2 create *,G state with CCEP as OIF.
- Both the peers send the *,G join towards RP and both the peers pull the traffic.


- When the traffic arrives, the S,G state is created on both the peers but only one of them will be forwarding based on the existing algorithm to determine which one of the MCT peers will keep the CCEP, OIF in forwarding and which one will keep it in MCT-blocked state.

MCT peer as intermediate Downstream router

P1 and P2 are the MCT peers and are acting as downstream routers for R1. R1 is the intermediate router.

P1, P2 and R1 are configured with PIM on the MCT VE interface. RP and source are beyond R1.

FIGURE 20 MCT peer as immediate downstream router



Hello exchange and neighbor state:

It acts and works as the upstream router.

Join or prune exchange and mcache state:

- The *,G joins come from the core to P2.
- P2 creates *,G state with uplink as OIF by consuming the join state.
- P2 due to its *,G state originates a join towards RP. This join is flooded on the MCT VLAN and R1 creates *,G state.
- P1 on receiving this join natively via ICL creates *,G state and adds ICL as OIF. Note that as a special case P1 will not include the *,G in the join it generates towards RP as in this case the IIF is CCEP and ICL is the only OIF and the remote CCEP is up. This is to avoid P1 pulling traffic from P2 unnecessarily on the ICL link because of P1 sending joins flooded on the VLAN and in turn P2 adds ICL as an OIF.

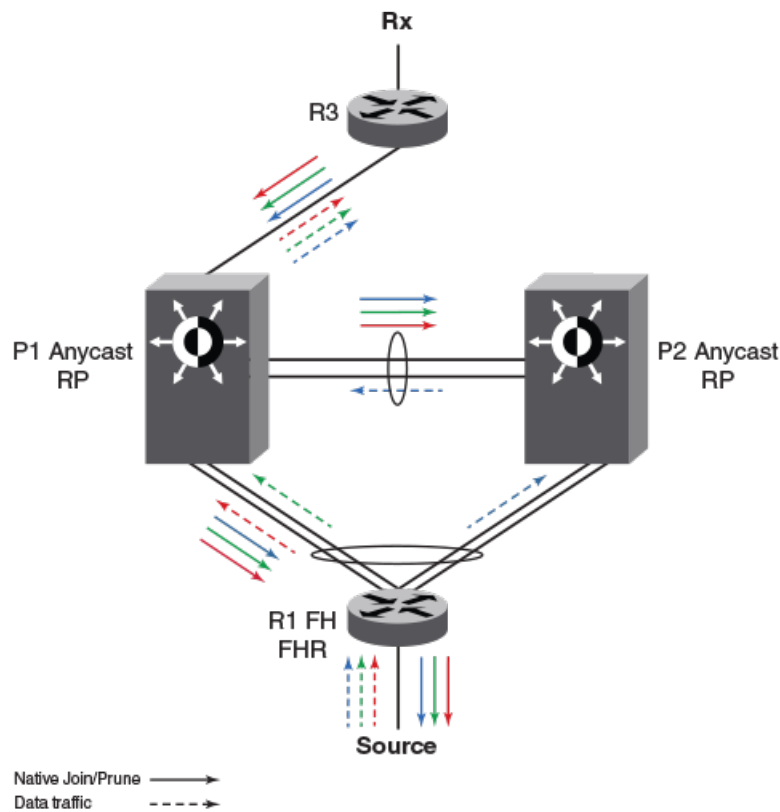
- R1 sends the join toward RP and pulls the traffic. Because the OIF at R1 is LAG, traffic pulled by R1 will be load-shared among the member links.
- Thus traffic for S,G will reach only one of the MCT peers. Assuming the traffic reaches P2, (S,G) state will be created on P2 and P2 will be forwarding the traffic.
- Assuming the traffic reaches P1 the traffic will be forwarded via ICL to P2 and P2 will forward it to its OIFs which is the link connecting to the core.

MCT peers as PIM Anycast RP

P1 and P2 are MCT peers and will be configured as Anycast RPs. This provides RP redundancy. In the event of one MCT peer going down, other will take over.

P1,P2 and R1 will be configured with PIM on the MCT VE interface. R1 is the FHR and receivers are connected to R3 which is present in the core.

FIGURE 21 MCT peer as PIM Anycast RP



Hello exchange and neighbor state:

This will have the same behavior as Upstream or Downstream router.

Join or prune exchange and mcache state:

- The receivers are connected to R3, it creates *,G state.
- R3 sends *,G join towards RP. As P1 is the anycast RP. P1 creates *,G state.


- When the multicast traffic starts, R1 sends registers to the configured RP address. Let's assume P2 receives the registers. In this case P2 will send register messages to P1. These register messages will not be sent via MDUP but will be sent as native register messages.
- P1 creates S,G state and sends S,G join towards R1, the join is sent on both ICL and CCEP ports.
- P2 processes this native join and creates S,G state adds ICL as immediate OIF, however doesn't send S,G join as the ICL is the only OIF and the CCEP is the IIF.
- If the native multicast data traffic comes to P2, P2 forwards it to P1 via ICL.
- If the native multicast data traffic comes to P1 it forwards it to R3.
- There are no S,G RPT prunes in this case.

Source directly connected to CEP on MCT VLAN

P1 and P2 are the MCT peers and R1 is the MCT Client. C1 and C2 are PIM routers connected to CEP ports.

P1, P2, R1, C1 and C2 will be configured with PIM on the MCT VE interface. RP is in the core and multicast source is connected to CEP router C1.

FIGURE 22 Source directly connected to CEP routers on MCT VLAN



Hello exchange and neighbor state:

P1, P2 and R1 will see each other as PIM neighbor. For C1 and C2 the following will happen.

- C1 originated Hello will be flooded on the MCT VLAN by P1. Thus P1, P2 and R1 will learn about C1.
- C1's Hello received by P2 on ICL will be flooded on MCT VLAN ports including the CEP ports where C2 is connected (sending it on local CCEP ports depends on the remote CCEP port status). Thus C2 creates neighbor state for C1. Similarly C1 creates neighbor state for C2.

- The Hellos from R1 received by either of the MCT peers will be sent via MDUP to the remote MCT peer.
- This Hello received via MDUP will be sent on CEP ports or other local CCEP ports.
- Thus both C1 and C2 learn about and create neighbor state for R1.

Source is connected to C1 on MCT VLAN:


- On receiving multicast traffic, C1 sends L2 registers to P1, P2, R1 and C2.
- Streams that have OIF on P1 update IIF as CEP and C1 as upstream neighbor and pulls traffic natively.
- Streams that have OIF on P2 update IIF as ICL and C1 as upstream neighbor and pulls traffic natively.
- Streams that have OIF on R1 update IIF as CCEP and C1 as upstream neighbor.

Multi tier MCT

P1 and P2 are MCT peers and P3 and P4 from another set of MCT peers. R1 and R2 are the clients for P1-P2 and P3-P4 clusters respectively. P1-P2 and P3-P4 are connected in Multi-Tier topology.

The source is connected to R2 on a VLAN other than the MCT VLAN. The receivers are connected to R1 on a VLAN other than the MCT VLAN.

FIGURE 23 Multi tier MCT



Hello exchange and neighbor state:

- Let's assume R1 sends hello to P2.
- P2 in turn sends it to CEP ports and the CCEP port towards peers P3 and P4. P2 also sends it to P1 via MDUP.
- P1 on receiving the hello via MDUP will send it on CEP ports and CCEP ports towards P3 and P4, if the remote CCEP ports towards P3 and P4 are down.
- Let's assume this Hello reaches P4 on its CCEP ports towards P1 and P2. P4 will send it on CCEP port towards R2 and will also send it to P3 via MDUP.
- Thus R2 learns about R1 and creates neighbor state. R1 also learns about R2 in a similar fashion.

Join or prune exchange and mcache state:

- The receivers are connected to R1, R1 creates *,G state and sends join towards RP.
- Let's assume P2 receives it. P2 then sends it to P1 via MDUP. Thus the MCT peers P1 and P2 both create *,G state.
- On receiving data traffic, R2 sends registers to the RP. Suppose it reaches P2. P2 then sends registers to P1.
- Both create S,G state and would send S,G join towards Source.
- Either P3 or P4 or both could receive the S,G joins. The MCT peer receiving it on its CCEP port will send it to the other peer via MDUP. Both P3 and P4 will process the received joins and create S,G state.
- Both Send S,G join towards R2 and data traffic from R2 could reach either P3 or P4 and will get forwarded load-shared on their CCEP towards P1 and P2.
- Streams could arrive at P1 and P2 (few at P1 and few at R2) and the streams which are received at the other peer will be pulled via ICL as the S,G joins are flooded on the MCT VLAN when it is sent towards P4 or P3.

Limitations

These are the limitations for MCT peers to support intermediate router functionality. These limitations are due to load-sharing and fast convergence trade-offs.

- PIM-DM is not supported.
- HLOS and Hitless Fail-over are not supported.
- Few packets may be lost during convergence interval or forwarding duplication may happen.
- MCT client will do flow based load-sharing, not per packet load-sharing.
- Traffic loss or duplication will happen when Keep-Alive VLAN, Cluster Communication Protocol (CCP) ,or ICL between MCT peers are not up.
- Both MCT peers must be configured as Anycast RPs for RP functionality. Configuring only one of the MCT peers as RP is not supported.
- Multicast routing configurations on session VLAN are not supported and are restricted in configuration. RP configured on the session VLAN's VE IP address is not supported.
- The load will only be shared, and may or may not be balanced across the CCEPs.
- During the convergence interval, a few packets may be lost. In the case of recoveries, some packets may end up being forwarded by both cluster routers during interval.
- Both the MCT peers maintain state and pull down traffic for all multicast flows from the core, whether the chassis is forwarding this stream to the local CCEP or not. This could potentially waste the bandwidth inside the core and on uplink.
- Administrators are expected to make the same static IGMP configurations on both the MCT peers for the CCEP. Mis-configuration is not identified.

Enabling PIM over MCT scaling optimization

By default, PIM over MCT has a scaling limit of 2K mcache (S,G) entries and 512 IGMP entries. You can increase the maximum scaling of these entries to 16K mcache entries and 4K IGMP entries.

The increase in scaling optimizes the MDUP between the MCT peers to minimize the processing of packets received from the CCEP ports downstream. The MDUP from one MCT peer to another peer occurs when one of the MCT peer sends the General Query (GS), Group Specific Query (GSQ), and Group Source Specific Query (GSSQ) queries. This peer is the querier in the MCT VLAN.

NOTE

This feature is not supported for snooping over MCT.

Perform the following steps to enable scaling optimization on both MCT peers.

1. In privileged EXEC mode, enter global configuration mode.

```
device# configure terminal
```

2. Enable IPv4 PIM over MCT scaling optimization.

```
device(config)# ip multicast-routing optimization mct-scaling
```

IPv6 MCT scaling optimization is also supported by the **ipv6 multicast-routing optimization mct-scaling** command.

3. Verify that the feature is enabled.

```
device(config)# show ip pim global
Global IPv4 PIM Settings
...
MCT Scaling Optimization : enabled
```

The following example is the configuration of the previous steps.

```
device# configure terminal
device(config)# ip multicast-routing optimization mct-scaling
```

Displaying IGMP and MLD cluster group information

To display the IGMP cluster groups, use the **show ip igmp cluster-client group** command.

```
device# show ip igmp cluster-client group
Total 1 groups
-----
Idx  Group Address      Port  Intf  GrpCmpV Mode  Timer Refreshed MDUPReq Srcs
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
  1  10.1.1.1            e1/7  v10   Ver2 exclude  237      N      N      0
-----
Total number of groups 1
Groups having cluster clients 1
```

To display the MLD cluster group, use the **show ipv6 mld cluster-client group** command.

```
device# show ipv6 mld cluster-client group 123:::3
Total 1 groups
-----
Idx  Group Address      Port  Intf  GrpCmpV Mode  Timer Refreshed MDUPReq Srcs
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
  1  123:::3            e1/7  v10   Ver1 exclude  253      Y      N      0
-----
```

To display the number of IGMP cluster groups, use the **show ip igmp group count** command.

```
device# show ip igmp group count
Total IGMP groups : 4096
```

Displaying MCT PIM Counters

To display the statistics and error counters for the Multicast MDUP channel between the MCT peers, use the **show ip pim counter mct** command.

```
device#show ip pim count mct
Multicast MCT Statistics for IPv4 (UP):
Messages assembled into the send buffer : 11811
Messages processed out of the recv buffer: 0
Segments sent successfully to TCP      : 11762
Segments failed to be accepted by TCP  : 0
```

```

Segments assembled into the receive buffer : 0
Messages dropped because (size > 1500) : 0
Messages dropped because it won't fit into available space in send buffer : 0
Segments dropped because it won't fit into available space in receive buffer: 0
Received messages dropped because of cluster-id mismatch : 0
Received messages dropped because the peer was not recognized : 0
Received messages dropped because cluster not active : 0
Received messages dropped because MCT VLAN unrecognized : 0
Received messages dropped because of bad message type : 0
Received messages dropped because of bad checksum : 0
Received bytes skipped because of sync or checksum errors : 0
PIM Hello Messages sent : 0
PIM J/P Messages sent : 0
PIM Assert Messages sent : 0
PIM Unknown not sent : 0
PIM Hello Messages received : 0
PIM J/P Messages received : 0
PIM Assert Messages received : 0
PIM Unknown received & dropped : 0
IGMPv1 reports sent : 0
IGMPv2 reports sent : 0
IGMPv3 reports sent : 0
IGMP leaves sent : 0
IGMP queries sent : 0
IGMP unknown not sent : 0
IGMPv1 reports received : 0
IGMPv2 reports received : 0
IGMPv3 reports received : 0
IGMP leaves received : 0
IGMP queries received : 0
IGMP unknown received & dropped : 0
device#

```

To display the MCT IPv6 PIM counter, use the **show ipv6 pim counter mct** command.

```

device#show ipv6 pim count mct
Multicast MCT Statistics for IPv6 (UP):
Messages assembled into the send buffer : 279
Messages processed out of the recv buffer: 523
Segments sent successfully to TCP : 279
Segments failed to be accepted by TCP : 0
Segments assembled into the receive buffer : 293
Messages dropped because (size > 1500) : 0
Messages dropped because it won't fit into available space in send buffer : 0
Segments dropped because it won't fit into available space in receive buffer: 0
Received messages dropped because of cluster-id mismatch : 0
Received messages dropped because the peer was not recognized : 0
Received messages dropped because cluster not active : 0
Received messages dropped because MCT VLAN unrecognized : 0
Received messages dropped because of bad message type : 0
Received messages dropped because of bad checksum : 0
Received bytes skipped because of sync or checksum errors : 0
PIM Hello Messages sent : 0
PIM J/P Messages sent : 0
PIM Assert Messages sent : 0
PIM Unknown not sent : 0
PIM Hello Messages received : 0
PIM J/P Messages received : 0
PIM Assert Messages received : 0
PIM Unknown received & dropped : 0
MLDv1 reports sent : 0
MLDv2 reports sent : 0
MLD leaves sent : 0
MLD queries sent : 0
MLD unknown not sent : 0
MLDv1 reports received : 0
MLDv2 reports received : 0
MLD leaves received : 0
MLD queries received : 0
MLD unknown received & dropped : 0
device#

```

Configuring a static multicast route

Static multicast routes allow you to control the network path used by multicast traffic. Static multicast routes are especially useful when the unicast and multicast topologies of a network are different. You can avoid the need to make the topologies similar by instead configuring static multicast routes.

You can configure more than one static multicast route. Extreme always uses the most specific route that matches a multicast source address. Thus, if you want to configure a multicast static route for a specific multicast source and also configure another multicast static route for all other sources, you can configure two static routes as shown in the examples below.

To add static routes to multicast router A (refer to [Configuring a static multicast route within a VRF](#) on page 140), enter commands such as the following.

```
PIMRouterA(config)# ip mroute 207.95.10.0 255.255.255.0 ethernet 1/2 distance 1
PIMRouterA(config)# ip mroute 0.0.0.0 0.0.0.0 ethernet 2/3 distance 1
PIMRouterA(config)# write memory
```

Syntax: `[no] ip mroute ip-addr ethernet slot/portnum | ve num [distance num]`

Or

Syntax: `[no] ip mroute ip-addr rpf_address rpf-num`

The **ip-addr** command specifies the PIM source for the route.

NOTE

In IP multicasting, a route is handled in terms of its source, rather than its destination.

You can use the **ethernet** slot/portnum parameter to specify a physical port or the **ve** num parameter to specify a virtual interface.

NOTE

The **ethernet** slot/portnum parameter do not apply to PIM SM.

The **distance** num parameter sets the administrative distance for the route. When comparing multiple paths for a route, the Extreme device prefers the path with the lower administrative distance.

NOTE

Regardless of the administrative distances, the Extreme device always prefers directly connected routes over other routes.

The **rpf_addressrpf-num** parameter specifies an RPF number.

The example above configures two static multicast routes. The first route is for a specific source network, 207.95.10.0/24. If the Extreme receives multicast traffic for network 207.95.10.0/24, the traffic must arrive on port 1/2. The second route is for all other multicast traffic. Traffic from multicast sources other than 207.95.10.0/24 must arrive on port 2/3.

[Configuring a static multicast route within a VRF](#) on page 140 shows an example of an IP Multicast network. The two static routes configured in the example above apply to this network. The commands in the example above configure PIM router A to accept PIM packets from 207.95.10.0/24 when they use the path that arrives at port 1/2, and accept all other PIM packets only when they use the path that arrives at port 2/3.

The distance parameter sets the administrative distance. This parameter is used by the software to determine the best path for the route. Thus, to ensure that the Extreme uses the default static route, assign a low administrative distance value. When comparing multiple paths for a route, the Extreme prefers the path with the lower administrative distance.

Configuring a static multicast route within a VRF

The Multi-Service IronWare software allows you to configure a static multicast route within a virtual routing instance (VRF). The static multicast route is defined within the VRF configuration as shown in the following.

```
device(config)# vrf vpn1
device(config-vrf-vpn1)#address family ipv4
device(config-vrf-vpn1)# ip mroute 105.105.105.0/24 14.14.14.105
```

Syntax: [no] vrf *vrf-name*

Syntax: [no] ip mroute *ip-addr*

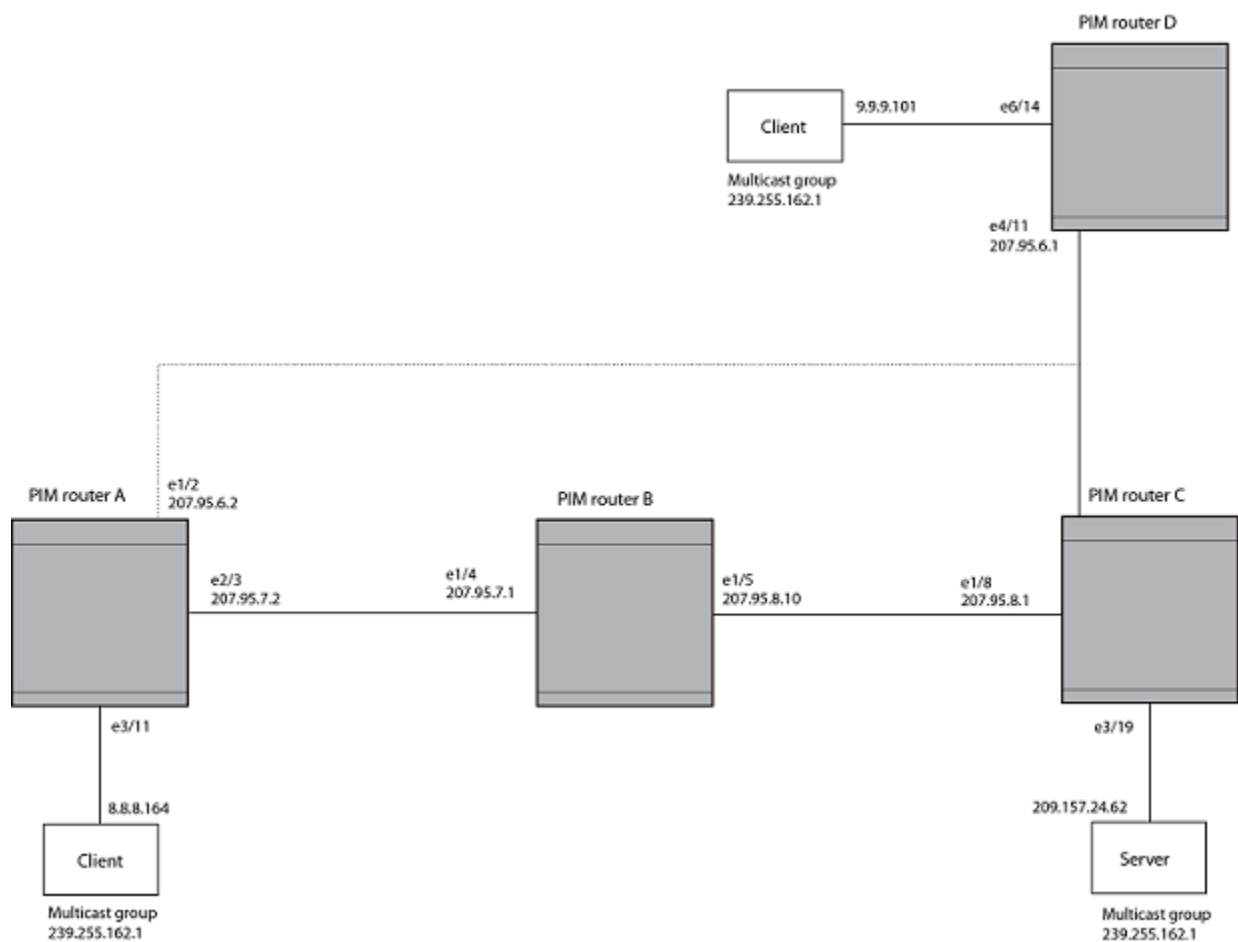
The **vrf** parameter specifies the virtual routing instance (VRF) specified by the variable **vrf-name** .

The **ip-addr** parameter specifies the destination IP address.

NOTE

Configuring a static multicast route for the default VRF is still accomplished using the command described in [Configuring a static multicast route](#) on page 139.

FIGURE 24 Example multicast static routes



To add a static route to a virtual interface, enter commands such as the following.

```
device(config)# ip mroute 0.0.0.0 0.0.0.0 ve 1 distance 1
device(config)# write memory
```

IGMP V3

The Internet Group Management Protocol (IGMP) allows an IPv4 system to communicate IP Multicast group membership information to its neighboring routers. The routers in turn limit the multicast of IP packets with multicast destination addresses to only those interfaces on the router that are identified as IP Multicast group members.

In IGMP V2, when a router sent a query to the interfaces, the clients on the interfaces respond with a membership report of multicast groups to the router. The router can then send traffic to these groups, regardless of the traffic source. When an interface no longer needs to receive traffic from a group, it sends a leave message to the router which in turn sends a group-specific query to that interface to see if any other clients on the same interface is still active.

In contrast, IGMP V3 provides selective filtering of traffic based on traffic source. A router running IGMP V3 sends queries to every multicast enabled interface at the specified interval. These general queries determine if any interface wants to receive traffic from the router. The following are the three variants of the Query message:

- A "General Query" is sent by a multicast router to learn the complete multicast reception state of the neighboring interfaces. In a General Query, both the Group Address field and the Number of Sources (N) field are zero.
- A "Group-Specific Query" is sent by a multicast router to learn the reception state, with respect to a "single" multicast address, of the neighboring interfaces. In a Group-Specific Query, the Group Address field contains the multicast address of interest, and the Number of Sources (N) field contains zero.
- A "Group-and-Source-Specific Query" is sent by a multicast router to learn if any neighboring interface desires reception of packets sent to a specified multicast address, from any of a specified list of sources. In a Group-and-Source-Specific Query, the Group Address field contains the multicast address of interest, and the Source Address [i] fields contain the source address(es) of interest.

The interfaces respond to these queries by sending a membership report that contains one or more of the following records that are associated with a specific group:

- Current-State Record that indicates from which sources the interface wants to receive and not receive traffic. The record contains source address of interfaces and whether or not traffic will be received or included (IS_IN) or not received or excluded (IS_EX) from that source.
- Filter-mode-change record. If the interface changes its current state from IS_IN to IS_EX, a TO_EX record is included in the membership report. Likewise, if an interface's current state changes from IS_EX to IS_IN, a TO_IN record appears in the membership report.

IGMP V2 Leave report is equivalent to a TO_IN (empty) record in IGMP V3. This record means that no traffic from this group will be received regardless of the source.

An IGMP V2 group report is equivalent to an IS_EX (empty) record in IGMP V3. This record means that all traffic from this group will be received regardless of source.

- Source-List-Change Record. If the interface wants to add or remove traffic sources from its membership report, the membership report can have an ALLOW record, which contains a list of new sources from which the interface wishes to receive traffic. It can also contains a BLOCK record, which lists current traffic sources from which the interfaces wants to stop receiving traffic.

In response to membership reports from the interfaces, the router sends a Group-Specific or a Group-and-Source Specific query to the multicast interfaces. For example, a router receives a membership report with a Source-List-Change record to block old sources from an

interface. The router sends Group-and-Source Specific Queries to the source and group (S,G) identified in the record. If none of the interfaces is interested in the (S,G), it is removed from (S,G) list for that interface on the router.

Each IGMP V3-enabled router maintains a record of the state of each group and each physical port within a virtual routing interface. This record contains the group, group-timer, filter mode, and source records information for the group or interface. Source records contain information on the source address of the packet and source timer. If the source timer expires when the state of the group or interface is in Include mode, the record is removed.

Default IGMP version

IGMP V3 is available for Extreme devices; however, these routers are shipped with IGMP V2-enabled. You must enable IGMP V3 globally or per interface.

Also, you can specify what version of IGMP you want to run on a device globally, on each interface (physical port or virtual routing interface), and on each physical port within a virtual routing interface. If you do not specify an IGMP version, IGMP V2 will be used.

Compatibility with IGMP V1 and V2

Different multicast groups, interfaces, and routers can run their own version of IGMP. Their version of IGMP is reflected in the membership reports that the interfaces send to the router. Routers and interfaces must be configured to recognize the version of IGMP you want them to process.

An interface or router sends the queries and reports that include its IGMP version specified on it. It may recognize a query or report that has a different version. For example, an interface running IGMP V2 can recognize IGMP V3 packets, but cannot process them. Also, a router running IGMP V3 can recognize and process IGMP V2 packet, but when that router sends queries to an IGMP V2 interface, the downgraded version is supported, not the upgraded version.

If an interface continuously receives queries from routers that are running versions of IGMP that are different from what is on the interface, the interface logs warning messages in the syslog every five minutes. Reports sent by interfaces to routers that contain different versions of IGMP do not trigger warning messages; however, you can see the versions of the packets using the **show ip igmp traffic** command.

The version of IGMP can be specified globally, per interface (physical port or virtual routing interface), and per physical port within a virtual routing interface. The IGMP version set on a physical port within a virtual routing interface supersedes the version set on a physical or virtual routing interface. Likewise, the version on a physical or virtual routing interface supersedes the version set globally on the device. The sections below present how to set the version of IGMP.

Globally enabling the IGMP version

To globally identify the IGMP version on an Extreme device, enter the following command.

```
device(config)# ip igmp version 3
```

Syntax: **[no] ip igmp version** *version-number*

Enter 1, 2, or 3 for Version-number. Version 2 is the default version.

Enabling the IGMP version per interface setting

To specify the IGMP version for a physical port, enter a command such as the following.

```
device(config)# interface eth 1/5
device(config-if-1/5)# ip igmp version 3
```

To specify the IGMP version for a virtual routing interface on a physical port, enter a command such as the following.

```
device(config)# interface ve 3
device(config-vif-1) ip igmp version 3
```

Syntax: `[no] ip igmp version version-number`

Enter 1, 2, or 3 for version-number . Version 2 is the default version.

Enabling the IGMP version on a physical port within a virtual routing interface

To specify the IGMP version recognized by a physical port that is a member of a virtual routing interface, enter a command such as the following.

```
device(config)# interface ve 3
device(config-vif-3)# ip igmp version 2
device(config-vif-3)# ip igmp port-version 3 e1/3 to e1/7 e2/9
```

In this example, the second line sets IGMP V2 on virtual routing interface 3. However, the third line set IGMP V3 on ports 1/3 through 1/7 and port e2/9. All other ports in this virtual routing interface are configured with IGMP V2.

Syntax: `[no] ip igmp port-version version-number ethernet port-number`

Enter 1, 2, or 3 for version-number . IGMP V2 is the default version.

The **ethernet** port-number parameter specifies which physical port within a virtual routing interface is being configured.

Enabling membership tracking and fast leave

NOTE

The IGMP V3 fast leave feature is supported in include mode, but does not work in the exclude mode.

IGMP V3 provides membership tracking and fast leave of clients. In IGMP V2, only one client on an interface needs to respond to a router's queries; therefore, some of the clients may be invisible to the router, making it impossible for the switch to track the membership of all clients in a group. Also, when a client leaves the group, the switch sends group specific queries to the interface to see if other clients on that interface need the data stream of the client who is leaving. If no client responds, the switch waits three seconds before it stops the traffic.

IGMP V3 contains the tracking and fast leave feature that you enable on virtual routing interfaces. Once enabled, all physical ports on that virtual routing interface will have the feature enabled. IGMP V3 requires all clients to respond to general and group specific queries so that all clients on an interface can be *tracked*. *Fast leave* allows clients to leave the group without the three second waiting period, if the following conditions are met:

- If the interface, to which the client belongs, has IGMP V3 clients only. Therefore, all physical ports on a virtual routing interface must have IGMP V3 enabled and no IGMP V1 or V2 clients can be on the interface. (Although IGMP V3 can handle V1 and V2 clients, these two clients cannot be on the interface in order for fast leave to take effect.)
- No other client on the interface is receiving traffic from the group to which the client belongs.

Every group on the physical interface of a virtual routing interface keeps its own tracking record. It can track by (source, group).

For example, two clients (Client A and Client B) belong to group1 but each is receiving traffic streams from different sources. Client A receives a stream from (source_1, group1) and Client B receives it from (source_2, group1). Now, if Client B leaves, the traffic stream (source_2, group1) will be stopped immediately. The **show ip igmp group tracking** command displays that clients in a group that are being tracked.

If a client sends a leave message, the client is immediately removed from the group. If a client does not send a report during the specified group membership time (the default is 140 seconds), that client is removed from the tracking list.

To enable the tracking and fast leave feature, enter commands such as the following.

```
device(config)# interface ve 13
device(config-vif-13)# ip igmp tracking
```

Syntax: **[no]** ip igmp tracking

Creating a static IGMP group

To configure a physical port to be a permanent (static) member of an IGMP group, enter the following commands.

```
device(config)# interface ethernet 1/5
device(config-if-e1000-1/5)# ip igmp static-group 224.10.1.1
```

Syntax: **[no]** ip igmp static-group *ip-address*

Enter the IP address of the static IGMP group for *ip-address*.

To configure a virtual port to be a permanent (static) member of an IGMP group, enter the following commands.

```
device(config)# interface ve 10
device(config-vif-10)# ip igmp static-group 224.10.1.1 ethernet 1/5
```

Syntax: **[no]** ip igmp *ip-address* static-group *ip-address* ethernet *slot-number/port-number*

Enter the IP address of the static IGMP group for *ip-address*.

Enter the ID of the physical port of the VLAN that will be a member of the group for **ethernet** *slot-number/port-number*.

NOTE

IGMPv3 does not support static IGMP group members.

NOTE

IGMPv3 is not supported for groups with L2 Multicast CPU protection enabled.

NOTE

Static IGMP groups are supported only in Layer 3 mode.

Setting the query interval

The IGMP query interval period defines how often a switch will query an interface for group membership. Possible values are 2-3600 seconds and the default value is 125 seconds, but the value you enter must be a little more than twice the group membership time.

To modify the default value for the IGMP query interval, enter the following.

```
device(config)# ip igmp query-interval 120
```

Syntax: **[no]** ip igmp query-interval *2-3600*

The interval must be a little more than two times the group membership time.

Setting the group membership time

Group membership time defines how long a group will remain active on an interface in the absence of a group report. Possible values are from 5 - 26000 seconds and the default value is 260 seconds.

To define an IGMP membership time of 240 seconds, enter the following.

```
device(config)# ip igmp group-membership-time 240
```

Syntax: [no] ip igmp group-membership-time 5-26000

Setting the maximum response time

The maximum response time defines the maximum number of seconds that a client can wait before it replies to the query sent by the router. Possible values are 1 - 25. The default is 10.

To change the IGMP maximum response time, enter a command such as the following at the global CONFIG level of the CLI.

```
device(config)# ip igmp max-response-time 8
```

Syntax: [no] ip igmp max-response-time *num*

The *num* parameter specifies the maximum number of seconds for the response time. Enter a value from 1 - 25. The default is 10.

Displaying IGMPv3 information

The sections below present the show commands available for IGMP V3.

Displaying IGMP group status

How to display the status of IGMP multicast groups on a device.

IGMP groups must be configured on the device.

The following steps can be performed in any order.

1. **show ip igmp group**

```
device# show ip igmp group
Total 293 groups
```

Idx	Group Address	Port	Intf	GrpCmpV	Mode	Timer	Srcs
1	225.0.0.1	e1/6	e1/6	Ver2	exclude	166	0
2	225.0.0.2	e1/6	e1/6	Ver2	exclude	166	0
3	225.0.0.3	e1/6	e1/6	Ver2	exclude	167	0
4	225.0.0.4	e1/6	e1/6	Ver2	exclude	167	0
5	225.0.0.5	e1/6	e1/6	Ver2	exclude	167	0
6	225.0.0.6	e1/6	e1/6	Ver2	exclude	168	0
7	225.0.0.7	e1/6	e1/6	Ver2	exclude	168	0
8	225.0.0.8	e1/6	e1/6	Ver2	exclude	168	0
9	225.0.0.9	e1/6	e1/6	Ver2	exclude	169	0
10	225.0.0.10	e1/6	e1/6	Ver2	exclude	169	0
11	225.0.0.11	e1/6	e1/6	Ver2	exclude	169	0
12	225.0.0.12	e1/6	e1/6	Ver2	exclude	170	0
13	225.0.0.13	e1/6	e1/6	Ver2	exclude	170	0
14	225.0.0.14	e1/6	e1/6	Ver2	exclude	170	0
15	225.0.0.15	e1/6	e1/6	Ver2	exclude	171	0
16	225.0.0.16	e1/6	e1/6	Ver2	exclude	171	0
17	225.0.0.17	e1/6	e1/6	Ver2	exclude	171	0
18	225.0.0.18	e1/6	e1/6	Ver2	exclude	172	0
19	225.0.0.19	e1/6	e1/6	Ver2	exclude	172	0

Displays the status of all IGMP multicast groups on a device.

2. `show ip igmp [vrf vrf-name] group group-address [detail]`

```
device# show ip igmp group 239.0.0.1 detail
Total 2 entries
-----
Idx Group Address Port Intf  GrpCmpV  Mode   Timer Srcs
-----+-----+-----+-----+-----+-----+-----
1   226.0.0.1   e6/2 v30   Ver2   exclude 218   2
S: 40.40.40.12
S: 40.40.40.11
S: 40.40.40.10
S: 40.40.40.2 (Age: 218)
S: 40.40.40.3 (Age: 218)
226.0.0.1 e6/3 e6/3 include 0 3
S: 30.30.30.3 (Age: 165)
S: 30.30.30.2 (Age: 165)
S: 30.30.30.1 (Age: 165)
```

Displays the status of one IGMP multicast group, in detail.

3. **NOTE**

If the tracking and fast leave feature is enabled, you can display the list of clients that belong to a particular group.

`show ip igmp [vrf vrf-name] group group-address [tracking]`

```
Brocade# show ip igmp group 224.1.10.1 tracking
Total 2 entries
-----
Idx Group Address Port Intf  GrpCmpV  Mode   Timer Srcs
-----+-----+-----+-----+-----+-----+-----
1   226.0.0.1   e6/2 v30   Ver2   exclude 253   3
S: 40.40.40.12
S: 40.40.40.11
S: 40.40.40.10
S: 40.40.40.2 (Age: 253)
C: 10.10.10.1 (Age: 253)
S: 40.40.40.3 (Age: 253)
C: 10.10.10.1 (Age: 253)
226.0.0.1 e6/3 e6/3 include 0 3
S: 30.30.30.3 (Age: 196)
C: 10.2.0.1 (Age: 196)
S: 30.30.30.2 (Age: 196)
C: 10.2.0.1 (Age: 196)
S: 30.30.30.1 (Age: 196)
C: 10.2.0.1 (Age: 196)
```

If you want a report for a specific multicast group, enter that group's address for *group-address*. Omit the *group-address* if you want a report for all multicast groups.

- The **vrf** parameter specifies that you want to display IGMP group information for the VRF specified by the *vrf-name* variable.
- Enter **detail** if you want to display the source list of the multicast group.
- Enter **tracking** if you want information on interfaces that have tracking enabled.

NOTE

IGMP v2 and v3 statistics are displayed on the report for each interface.

Displays the IGMP group information for the group address, 224.1.10.1. This example assumes that the tracking and fast leave option is enabled.

```
device# show ip igmp group 224.1.10.1 tracking

Total 2 entries
-----
Idx  Group Address      Port  Intf  GrpCmpV  Mode    Timer  Srcs
---  -
1    226.0.0.1          e6/2  v30   Ver2     exclude 253    3
   S: 40.40.40.12
   S: 40.40.40.11
   S: 40.40.40.10
   S: 40.40.40.2 (Age: 253)
   C: 10.10.10.1 (Age: 253)
   S: 40.40.40.3 (Age: 253)
   C: 10.10.10.1 (Age: 253)
   226.0.0.1          e6/3  e6/3  include  0       3
   S: 30.30.30.3 (Age: 196)
   C: 10.2.0.1 (Age: 196)
   S: 30.30.30.2 (Age: 196)
   C: 10.2.0.1 (Age: 196)
   S: 30.30.30.1 (Age: 196)
   C: 10.2.0.1 (Age: 196)
```

For an explanation of the fields shown in the example above, see the following table:

Field	Description
Group	The address of the multicast group.
Port	The physical port on which the multicast group was received.
Intf	The virtual interface on which the multicast group was received.
GrpCmpV	The version of the IGMP group.
Timer	Shows the number of seconds the interface can remain in exclude mode. An exclude mode changes to include mode if it does not receive an "IS_EX" or "TO_EX" message during a certain period of time. The default is 140 seconds.
Mode	Indicates current mode of the interface: include or exclude. If the interface is in Include mode, it admits traffic only from the source list. If an interface is in exclude mode, it denies traffic from the source list and accepts the rest.
Srcs	Identifies the source list that will be included or excluded on the interface. NOTE If the IGMP V2 group is in exclude mode with a #_src of 0, the group excludes traffic from 0 (zero) source list, which means that all traffic sources are included.

Clearing the IGMP group membership table

To clear the IGMP group membership table, enter the following command.

```
device# clear ip igmp cache
```

Syntax: `clear ip igmp [vrf vrf-name] cache`

This command clears the IGMP membership for the default router instance or for a specified VRF.

Use the **vrf** option to clear the traffic information for a VRF instance specified by the **vrf-name** variable.

Displaying static IGMP groups

The following command displays static IGMP groups for the "eng" VRF.

```
device#show ip igmp vrf eng static
Group Address      Interface Port List
-----+-----+-----
      229.1.0.12    4/1 ethe 4/1
      229.1.0.13    4/1 ethe 4/1
      229.1.0.14    4/1 ethe 4/1
      229.1.0.92    4/1 ethe 4/1
```

Syntax: `show ip igmp [vrf vrf-name] static`

The **vrf** parameter specifies that you want to display static IGMP group information for the VRF specified by the *vrf-name* variable.

TABLE 27 Output of show ip igmp vrf static

This field	Displays
Group Address	The address of the multicast group.
Interface Port List	The physical ports on which the multicast groups are received.

Displaying the IGMP status of an interface

You can display the status of a multicast enabled port by entering a command such as the following.

```
device# show ip igmp interface
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
Intf/Port|Groups| Version |Querier      | Timer  |V1Rtr|V2Rtr|Tracking
      |      |Oper  Cfg|              | |Qrr GenQ|      |      |
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
e6/3      | 1    | 3    3 Self  |      0  94 No  No  | Disabled
e6/4      | 0    | 2    - Self  |      0  94 No  No  | Disabled
v30       | 1    | 3    3      |      0  20 No  No  | Disabled
v40       | e6/2 | 3    - Self  |      0  20 No  No  | Disabled
v40       | e6/2 | 3    3      |      0  20 No  No  | Disabled
v50       | 0    | 2    -      |      0  29 No  No  | Disabled
v50       | e12/1| 2    - Self  |      0  29 No  No  | Disabled
v50       | e6/8 | 2    - 50.1.1.10 | 46  0 No  Yes | Disabled
v50       | e6/1 | 2    - Self  |      0 115 No  Yes | Disabled
```

Syntax: `show ip igmp [vrf vrf-name] interface [ve number | ethernet port-address | tunnel num]`

The **vrf** parameter specifies that you want to display IGMP interface information for the VRF specified by the *vrf-name* variable.

Enter **ve** and its *number*, or **ethernet** and its *port-address* to display information for a specific virtual routing interface, or ethernet interface.

The **tunnelnum** parameter specifies a GRE tunnel interface that is being configured. The GRE tunnel interface is enabled under the router PIM configuration.

Entering an address for *group-address* displays information for a specified group on the specified interface.

The report shows the following information:

TABLE 28 Output of show ip igmp interface

This field	Displays
Intf	The virtual interface on which IGMP is enabled.
Port	The physical port on which IGMP is enabled.
Groups	The number of groups that this interface or port has membership.

TABLE 28 Output of show ip igmp interface (continued)

This field	Displays
Version	
Oper	The IGMP version that is operating on the interface.
Cfg	The IGMP version that is configured for this interface.
Querier	Where the Querier resides: The IP address of the router where the querier is located or Self - if the querier is on the same router as the intf or port.
Max response	
oQrr	Other Querier present timer.
GenQ	General Query timer
V1Rtr	Whether IGMPv1 is present on the intf or port.
V2Rtr	Whether IGMPv2 is present on the intf or port.
Tracking	Fast tracking status: Enabled or Disabled

Displaying IGMP traffic status

To display the traffic status on each virtual routing interface, enter the following command.

```
device# show ip igmp traffic
Recv  QryV2  QryV3  G-Qry  GSQry  MbrV2  MbrV3  Leave  IsIN  IsEX  ToIN  ToEX  ALLOW  BLK
v5    29     0     0     0     0     0     0     0     0     0     0     0     0
v18   15     0     0     0     0     30    0     60    0     0     0     0     0
v110  0     0     0     0     0     97    0     142   37    2     2     3     2
Send  QryV1  QryV2  QryV3  G-Qry  GSQry
v5    0     2     0     0     0
v18   0     0     30    30    0
v110  0     0     30    44    11
```

Syntax: show ip igmp [vrf *vrf-name*] traffic

The **vrf** parameter specifies that you want to display IGMP traffic information for the VRF specified by the *vrf-name* variable.

The report shows the following information:

TABLE 29 Output of show ip igmp vrf traffic

This field	Displays
QryV2	Number of general IGMP V2 query received or sent by the virtual routing interface.
QryV3	Number of general IGMP V3 query received or sent by the virtual routing interface.
G-Qry	Number of group specific query received or sent by the virtual routing interface.
GSQry	Number of source specific query received or sent by the virtual routing interface.
MbrV2	The IGMP V2 membership report.
MbrV3	The IGMP V3 membership report.

TABLE 29 Output of show ip igmp vrf traffic (continued)

This field	Displays
Leave	Number of IGMP V2 "leave" messages on the interface. (See ToEx for IGMP V3.)
IsIN	Number of source addresses that were included in the traffic.
IsEX	Number of source addresses that were excluded in the traffic.
ToIN	Number of times the interface mode changed from exclude to include.
ToEX	Number of times the interface mode changed from include to exclude.
ALLOW	Number of times that additional source addresses were allowed or denied on the interface:
BLK	Number of times that sources were removed from an interface.

Clearing IGMP traffic statistics

To clear statistics for IGMP traffic, enter the following command.

```
device# clear ip igmp traffic
```

Syntax: `clear ip igmp [vrf vrf-name] traffic`

This command clears all the multicast traffic information on all interfaces on the device.

Use the **vrf** option to clear the traffic information for a VRF instance specified by the **vrf-name** variable. T

Displaying IGMP settings

To display global IGMP settings or IGMP settings for a specified VRF. To display global IGMP settings, enter the following command.

```
device show ip igmp settings
IGMP Global Configuration
  Query Interval      : 125s
  Configured Query Interval : 125s
  Max Response Time   : 10s
  Group Membership Time : 260s
  Configured Version  : 2
  Operating Version   : 2
```

Syntax: `show ip igmp [vrf vrf-name] settings`

The **vrf** parameter specifies that you want to display IGMP settings information for the VRF specified by the **vrf-name** variable.

The report shows the following information:

TABLE 30 show ip igmp output

This field	Displays
Query Interval	How often the router will query an interface for group membership.
Configured Query Interval	The query interval that has been configured for the router.
Max Response Time	The length of time in seconds that the router will wait for an IGMP (V1 or V2) response from an interface before concluding that the group member on that interface is down and removing it from the group.
Group Membership Time	The length of time in seconds that a group will remain active on an interface in the absence of a group report.
Configured Version	The IGMP version configured on the router.
Operating Version	The IGMP version operating on the router.

Source-specific multicast

Using the Any-Source Multicast (ASM) service model, sources and receivers register with a multicast address. The protocol uses regular messages to maintain a correctly configured broadcast network where all sources can send data to all receivers and all receivers get broadcasts from all sources.

With Source-specific multicast (SSM), the "channel" concept is introduced where a "channel" consists of a single source and multiple receivers who specifically register to get broadcasts from that source. Consequently, receivers are not burdened with receiving data they have no interest in, and network bandwidth requirements are reduced because the broadcast need only go to a sub-set of users. The address range 232/8 has been assigned by the Internet Assigned Numbers Authority (IANA) for use with SSM.

IGMP V3 and source specific multicast protocols

When IGMP V3 and PIM Sparse (PIM-SM) is enabled, the source specific multicast service (SSM) can be configured. SSM simplifies PIM-SM by eliminating the RP and all protocols related to the RP. IGMPv3 and PIM-SM must be enabled on any ports that you want SSM to operate.

Configuring PIM SSM group range

PIM Source Specific Multicast (SSM) is a subset of the PIM SM protocol. In PIM SSM mode, the shortest path tree (STP) is created at the source. The STP is created between the receiver and source, but the STP is built without the help of the RP. The router closest to the interested receiver host is notified of the unicast IP address of the source for the multicast traffic. PIM SSM goes directly to the source-based distribution tree without the need of the RP connection. PIM SSM is different from PIM SM because it forms its own STP tree, without forming a shared tree. The multicast address group range is 232.0.0.0/8.

To configure a single SSM group address, enter the following command under the router pim configuration:

```
device(config)#router pim
device(config-pim-router)#ssm-enable range 232.1.1.1/8
```

Syntax: [no] ssm-enable range *group-address address-mask*

The **group-address** parameter specifies the multicast address for the SSM address range. If this is not configured, the range will default to 232/8 as assigned by the Internet Assigned Numbers Authority (IANA) for use with SSM.

The **address-mask** parameter specifies the mask for the SSM address range.

To disable SSM, use the [no] form of this command.

Displaying source-specific multicast configuration information

To display PIM Sparse configuration information, use the **show ip pim sparse** command as described in [Displaying basic PIM Sparse configuration information](#) on page 91.

Configuring multiple SSM group ranges

The **ssm-enable range acl-id/acl-name** command allows you to configure multiple SSM group ranges using an ACL.

Configuration Considerations

- The existing **ssm-enable range group-address address-mask** command will continue to exist.

- The ACL must be configured with the SSM group address in the permit clause of the **ssm-enable range** *acl-id* or *acl-name* command. If the **ssm-enable range group-address address-mask** command permits a clause, then that group will also operate in the PIM-SM mode.
- If the **ssm-enable range acl-id** or **acl-name** command is configured with a non-existent or empty ACL, then the SSM group will operate in PIM-SM mode (non PIM-SSM mode). However when an ACL is added or updated, then the group will exist in a PIM-SSM mode. By default, an empty ACL will deny all.
- By default, the group address mentioned in the IGMPv2 ssm-mapping ACL will decide if the group address is a PIM-SSM group or non PIM-SSM group. Therefore, if a user wants to prevent a group from operating in PIM-SSM mode, then the user's configuration must consistently deny the group in all configuration options for PIM-SSM range.
- ACL of any type (named or unnamed, standard or extended) can be used to specify the SSM group range. If an extended ACL is used, then the destination ip address should be used to specify the group address. Any configuration in the source address of an extended ACL is ignored. Only permit statements are considered in the ACL configuration. Any deny statements in the ACL clause are also ignored.

To configure multiple SSM group address using an ACL, enter the following command under the router pim configuration:

```
device(config)#router pim
device(config-pim-router)#ssm-enable range xyz
```

The example displayed above configures PIM so that it uses the group addresses allowed by ACL, xyz as its PIM SSM range.

Syntax: **[no] ssm-enable range** *acl-id* or *acl-name*

The **acl-id/acl-name** parameter specifies the ACL id or name used to configure multiple SSM group ranges.

To disable the SSM mapping range ACL, use the **no** form of this command.

NOTE

The **ssm-enable range acl-id** or **acl-name** command also supports IPv6 traffic. The **ssm-enable range acl-id** or **acl-name** command must be configured under the IPv6 router pim configuration to support IPv6.

Displaying information for PIM SSM range ACL

To display information for PIM SSM range ACL configuration enter the following command at any CLI level:

```
deviceshow ip pim sparse
Global PIM Sparse Mode Settings
Maximum Mcache          : 0           Current Count           : 0
Hello interval          : 30           Neighbor timeout        : 105
Join/Prune interval     : 60           Inactivity interval     : 180
Register Suppress Time  : 60           Register Probe Time     : 10
SPT Threshold           : 1           Hardware Drop Enabled   : Yes
Bootstrap Msg interval : 60           Candidate-RP Msg interval : 60
Register Stop Delay     : 60           Register Suppress interval : 60
SSM Enabled             : Yes
SSM Group Range         : 224.1.1.1/24
SSM Group Range ACL     : xyz
Route Precedence        : mc-non-default mc-default uc-non-default uc-default
```

NOTE

The **show ipv6 pim sparse** command also displays PIM SSM range ACL configuration.

IGMPv2 SSM mapping

The PIM-SSM feature requires all IGMP hosts to send IGMPv3 reports. Where you have an IGMPv2 host, this can create a compatibility problem. In particular, the reports from an IGMPv2 host contain a Group Multicast Address but do not contain source addresses. The IGMPv3 reports contain both the Group Multicast Address and one or more source addresses. This feature converts IGMPv2 reports into IGMPv3 reports through use of the **ip igmp ssm-map** commands and a properly configured ACL.

The ACL used with this feature filters for the Group Multicast Address. The ACL is then associated with one or more source addresses using the **ip igmp ssm-map static** command. When the **ip igmp ssm-map enable** command is configured, IGMPv3 reports are sent for IGMPv2 hosts.

The following sections describe how to configure the ACL and the **ip igmp ssm-map** commands to use the IGMPv2 SSM mapping feature:

- Configuring an ACL for IGMPv2 SSM mapping
- Configuring the IGMPv2 SSM Mapping Commands

NOTE

IGMPv2 SSM Mapping is not supported for IGMP static groups.

Configuring an ACL for IGMPv2 SSM mapping

You can use either a standard or extended ACL to identify the group multicast address you want to add source addresses to when creating a IGMPv3 report.

For standard ACLs, you must create an ACL with a permit clause and the **ip-source-address** variable must contain the group multicast address. This can be configured directly with a subnet mask or with the **host** keyword in which case a subnet mask of all zeros (0.0.0.0) is implied.

In the following example, **access-list 20** is configured for the group multicast address: 224.1.1.0 with a subnet mask of 0.0.0.255.

```
device(config)# access-list 20 permit 224.1.1.0 0.0.0.255
```

In the following example, **access-list 20** is configured for the group multicast address: 239.1.1.1 by including the **host** keyword.

```
device(config)# access-list 20 host 239.1.1.1
```

For extended ACLs, the **source address** variable must contain either **000** or the **any** keyword. Additionally, the extended ACL must be configured with a **permit** clause and the **host** keyword. This can be configured directly with a subnet mask or with the **host** keyword in which case a subnet mask of all zeros (0.0.0.0) is implied.

The **ip-destination-address** variable must contain the group multicast address.

In the following example, **access-list 100** is configured for the group multicast address: 232.1.1.1 with a subnet mask of 0.0.0.255.

```
device(config)# access-list 20 permit 224.1.1.0 0.0.0.255
```

In the following example, **access-list 100** is configured for the group multicast address: 232.1.1.1.

```
device(config)# access-list 100 permit any host 232.1.1.1
```

Configuring the IGMPv2 SSM mapping commands

The **ip ssm-map** commands are used to enable the IGMPv2 mapping feature and to define the maps between IGMPv2 Group addresses and multicast source addresses as described in the following sections.

Enabling IGMPv2 SSM mapping

To enable the IGMPv2 mapping feature enter the command as shown in the following.

```
device(config)# ip igmp ssm-map enable
```

Syntax: `[no] ip igmp ssm-map enable`

The **no** option is used to turn off the IGMPv2 mapping feature that has previously been enabled.

Configuring the map between a IGMPv2 group address and a multicast source

To configure a map between an IGMPv2 Group address and a multicast source address use the **ip igmp ssm-map** command, as shown in the following.

```
device(config)# ip igmp ssm-map static 1.1.1.1
device(config)# ip igmp ssm-map 20 1.1.1.1
```

Syntax: `[no] ip igmp ssm-map acl-number-name source-address`

The `acl-number-name` variable specifies the ACL that contains the group multicast address.

The `source-address` variable specifies the source address that you want to map to the group multicast address specified in the ACL.

The **no** option is used to delete a previously configured SSM map.

Example configuration

In the following example configuration, one extended ACL and two standard ACLs are defined with group multicast addresses. The **ip igmp ssm-map** commands are configured to map the ACLs to source addresses and to enable the feature on the router.

```
device(config)# access-list 20 host 239.1.1.1
device(config)# access-list 20 permit 224.1.1.0 0.0.0.225
device(config)# access-list 100 permit any host 232.1.1.1
device(config)# ip igmp ssm-map static 20 1.1.1.1
device(config)# ip igmp ssm-map static 20 2.2.2.2
device(config)# ip igmp ssm-map static 100 1.1.1.1
device(config)# ip igmp ssm-map enable
```

Displaying an IGMP SSM mapping information

The **show ip igmp ssm-map** command displays the association between a configured ACL and source address mapped to it, as shown in the following.

```
device show ip igmp ssm-map
+-----+-----+
| Acl id | Source Address |
+-----+-----+
|      20 |      1.1.1.1   |
|     100 |      1.1.1.1   |
|      20 |      2.2.2.2   |
|      20 |      2.2.2.3   |
|      20 |      2.2.2.4   |
|      20 |      2.2.2.5   |
|      20 |      2.2.2.6   |
```

Syntax: `show ip igmp ssm-map`

The **show ip igmp ssm-map group-address** displays the ACL ID that has the specified multicast group address in its permit list and lists the source addresses mapped to the specified multicast group address, as shown in the following.

```
device show ip igmp ssm-map 232.1.1.1
+-----+-----+
```

Acl id	Source Address
20	1.1.1.1
100	1.1.1.1
20	2.2.2.2
20	2.2.2.3
20	2.2.2.4
20	2.2.2.5
20	2.2.2.6

Syntax: `show ip igmp ssm-map group-address`

IPv6 Multicast Routing


- IPv6 PIM Sparse 157
- PIM Anycast RP..... 183
- Multicast Listener Discovery and source-specific multicast protocols..... 186

IPv6 PIM Sparse

IPv6 Protocol Independent Multicast (PIM) Sparse is supported. IPv6 PIM Sparse provides multicasting that is especially suitable for widely distributed multicast environments.

In an IPv6 PIM Sparse network, an IPv6 PIM Sparse router that is connected to a host that wants to receive information for a multicast group must explicitly send a join request on behalf of the receiver (host).

FIGURE 25 Example IPv6 PIM Sparse domain



PIM Sparse router types

Routers that are configured with PIM Sparse interfaces also can be configured to fill one or more of the following roles:

- BSR - The Bootstrap Router (BSR) distributes RP information to the other PIM Sparse routers within the domain. Each PIM Sparse domain has one active BSR. For redundancy, you can configure ports on multiple routers as candidate BSRs. The PIM Sparse protocol uses an election process to select one of the candidate BSRs as the BSR for the domain. The BSR with the highest BSR priority (a user-configurable parameter) is elected. If the priorities result in a tie, then the candidate BSR interface with the highest IP address is elected. In the example in [IPv6 PIM Sparse](#) on page 157, PIM Sparse router B is the BSR. Port 2/2 is configured as a candidate BSR.
- RP - The Rendezvous Points (RP) is the meeting point for PIM Sparse sources and receivers. A PIM Sparse domain can have multiple RPs, but each PIM Sparse multicast group address can have only one active RP. PIM Sparse routers learn the addresses of RPs and the groups for which they are responsible from messages that the BSR sends to each of the PIM Sparse routers. In the example in [IPv6 PIM Sparse](#) on page 157, PIM Sparse router B is the RP. Port 2/2 is configured as a candidate Rendezvous Point (RP).

To enhance overall network performance, the device uses the RP to forward only the first packet from a group source to the group receivers. After the first packet, the device calculates the shortest path between the receiver and the source (the Shortest Path Tree, or SPT) and uses the SPT for subsequent packets from the source to the receiver. The device calculates a separate SPT for each source-receiver pair.

NOTE

It is recommended that you configure the same ports as candidate BSRs and RPs.

RP paths and SPT paths

[IPv6 PIM Sparse](#) on page 157 shows two paths for packets from the source for group fec0:1111::1 and a receiver for the group. The source is attached to PIM Sparse router A and the recipient is attached to PIM Sparse router C. PIM Sparse router B is the RP for this multicast group. As a result, the default path for packets from the source to the receiver is through the RP. However, the path through the RP sometimes is not the shortest path. In this case, the shortest path between the source and the receiver is over the direct link between router A and router C, which bypasses the RP (router B).

To optimize PIM traffic, the protocol contains a mechanism for calculating the Shortest Path Tree (SPT) between a given source and a receiver. PIM Sparse routers can use the SPT as an alternative to using the RP for forwarding traffic from a source to a receiver. By default, the device forwards the first packet it receives from a given source to a given receiver using the RP path, but subsequent packets from that source to that receiver through the SPT. In [IPv6 PIM Sparse](#) on page 157, router A forwards the first packet from group fec0:1111::1 source to the destination by sending the packet to router B, which is the RP. Router B then sends the packet to router C. For the second and all future packets that router A receives from the source for the receiver, router A forwards them directly to router C using the SPT path.

RFC 3513 and RFC 4007 compliance for IPv6 multicast scope-based forwarding

The IPv6 multicast implementation recognizes scopes and conforms to the scope definitions in RFC 3513. Per RFC 3513, scopes 0 and 3 are reserved and packets are not forwarded with an IPv6 destination multicast address of scopes 0 and 3. Additionally, scopes 1 and 2 are defined as Node-Local and Link-Local and are not forwarded. Thus, the implementation forwards only those packets with an IPv6 multicast destination address with scope 4 or higher.

RFC 4007 defines 'scope zones' and requires that the forwarding of packets received on any interface of a particular scope zone be restricted to that scope zone. Currently, the device supports one zone for each scope, and the default zone for scope 4 and higher consists of all interfaces in the system. Thus, the default zones for scope 4 and higher are the same size.

Configuring PIM Sparse

To configure the device for IPv6 PIM Sparse, perform the following tasks:

- Enable the IPv6 PIM Sparse of multicast routing.
- Configure an IPv6 address on the interface.
- Enable IPv6 PIM Sparse.
- Identify the interface as an IPv6 PIM Sparse border, if applicable.
- Enable IPv6 Protocol Independent Multicast Sparse mode (PIM-SM) for a specified VRF, if applicable.
- Identify the device as a candidate PIM Sparse Bootstrap Router (BSR), if applicable.
- Identify the device as a candidate PIM Sparse Rendezvous Point (RP), if applicable.
- Specify the IP address of the RP (if you want to statically select the RP).

NOTE

It is recommended that you configure the same device as both the BSR and the RP.

IPv6 PIM-Sparse mode

To configure a device for IPv6 PIM Sparse, perform the following tasks:

- Identify the Layer 3 switch as a candidate sparse Rendezvous Point (RP), if applicable.
- Specify the IPv6 address of the RP (to configure statically).

The following example enables IPv6 PIM-SM routing. Enter the following command at the configuration level to enable IPv6 PIM-SM globally.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)#
```

To enable IPv6 PIM Sparse mode on an interface, enter commands such as the following.

```
device(config)# interface ethernet 2/2
device(config-if-e10000-2/2)# ipv6 address a000:1111::1/64
device(config-if-e10000-2/2)# ipv6 pim-sparse
```

Syntax: [no] ipv6 pim-sparse

Use the **no** option to remove IPv6 PIM sparse configuration from the interface.

The commands in this example add an IPv6 interface to port 2/2, then enable IPv6 PIM Sparse on the interface.

Configuring IPv6 PIM-SM on a virtual routing interface

You can enable IPv6 PIM-SM on a virtual routing interface by entering commands such as the following.

```
device(config)# interface ve 15
device(config-vif-15)# ipv6 address a000:1111::1/64
device(config-vif-15)# ipv6 pim-sparse
```

Enabling IPv6 PIM-SM for a specified VRF

To enable IPv6 PIM-SM for the VRF named "blue", create the VRF named "blue", enable it for IPv6 routing, and then enable IPv6 PIM-SM for the VRF, as shown in the following example.

```
device(config)# vrf blue
device(config-vrf-blue)# rd 11:1
device(config-vrf-blue)# address-family ipv6
device(config-vrf-blue-ipv6)# router pim
device(config-pim-router)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)
```

Syntax: `[no] ipv6 router pim [vrf vrf-name]`

The *vrf* parameter allows you to configure IPv6 PIM-SM on the virtual routing instance (VRF) specified by the *vrf-name* variable. All PIM parameters available for the default router instance are configurable for a VRF-based PIM instance.

Use the *no* option to remove all configuration for PIM multicast on the specified VRF.

Configuring BSRs

In addition to the global and interface parameters configured in the prior sections, you must identify an interface on at least one device as a candidate PIM Sparse Bootstrap Router (BSR) and a candidate PIM Sparse Rendezvous Point (RP).

NOTE

It is possible to configure the device as only a candidate BSR or an RP, but it is recommended that you configure the same interface on the same device as both a BSR and an RP.

To configure the device as a candidate BSR, enter commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# bsr-candidate ethernet 1/3 32 64
BSR address: 31::207, hash mask length: 32, priority: 64
```

This command configures Ethernet interface 1/3 as the BSR candidate with a mask length of 32 and a priority of 64.

To configure the device as a candidate BSR for a specified VRF, enter the commands as shown in the following example.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# bsr-candidate ethernet 1/3 32 64
BSR address: 31::207, hash mask length: 32, priority: 64
```

Syntax: `[no] bsr-candidate ethernet slot/portnum | loopback num | ve num hash-mask-length [priority]`

Use the *no* option to remove the candidate BSR configuration for a specified VRF.

The **ethernet slot/portnum | loopback num | ve num** parameter specifies the interface. The device will advertise the specified interface's IP address as a candidate BSR:

- Enter **ethernet slot/portnum** for a physical interface (port).
- Enter **loopback num** for a loopback interface.
- Enter **ve num** for a virtual interface.

The *hash-mask-length* variable specifies the number of bits in a group address that are significant when calculating the group-to-RP mapping. You can specify a value from 1 through 32.

The *priority* variable specifies the BSR priority. You can specify a value from 0 through 255. When the election process for BSR takes place, the candidate BSR with the highest priority becomes the BSR. The default is 0.

Setting the BSR message interval

The BSR message interval timer defines the interval at which the BSR sends RP candidate data to all IPv6-enabled routers within the IPv6 PIM Sparse domain. The default is 60 seconds.

To set the IPv6 PIM BSR message interval timer to 16 seconds, enter commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# bsr-msg-interval 16
Changed BSR message interval to 16 seconds.
```

To set the IPv6 PIM BSR message interval timer to 16 seconds for a specified VRF, enter the commands as shown in the following example.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# bsr-msg-interval 16
Changed BSR message interval to 16 seconds.
```

Syntax: `[no] bsr-msg-interval num`

The *num* parameter specifies the number of seconds and can be from 10 - 65535. The default is 60.

Use the *no* option to disable a timer that has been configured.

Configuring candidate RP

Enter a command such as the following to configure the device as a candidate RP.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# rp-candidate ethernet 2/2
```

To configure the device as a candidate RP for a specified VRF, enter the commands as shown in the following example.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# rp-candidate ethernet 2/2
```

Syntax: `[no] rp-candidate ethernet slot/portnum | loopback num | ve num`

The **ethernet** *slot/portnum* | **loopback** *num* | **ve** *num* parameter specifies the interface. The device will advertise the specified interface IP address as a candidate RP:

- Enter **ethernet** *slot/portnum* for a physical interface (port).
- Enter **loopback** *num* for a loopback interface.
- Enter **ve** *num* for a virtual interface.

To add address ranges for which the device is a candidate RP, enter commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# rp-candidate add ff02::200:2 64
```

To add address ranges for a specified VRF for which the device is a candidate RP, enter commands such as the following.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# rp-candidate add ff02::200:2 64
```

Syntax: `[no] rp-candidate add group-ipv6 address mask-bits`

You can delete the configured RP candidate group ranges by entering commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# rp-candidate delete ff02::200:1 128
```

You can delete the configured RP candidate group ranges for a specified VRF by entering commands such as the following:

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router-vrf-blue)# rp-candidate delete ff02::200:1 128
```

Syntax: `[no] rp-candidate delete group-ipv6 address mask-bits`

The usage for the `group-ipv6 addressmask-bits` parameter is the same as for the `rp-candidate add` command.

Statically specifying the RP

It is recommended that you use the IPv6 PIM Sparse mode RP election process so that a backup RP can automatically take over if the active RP router becomes unavailable. However, if you do not want the RP to be selected by the RP election process but instead you want to explicitly identify the RP by its IPv6 address, use the `rp-address` command.

If you explicitly specify the RP, the device uses the specified RP for all group-to-RP mappings and overrides the set of candidate RPs supplied by the BSR.

NOTE

Specify the same IP address as the RP on all IPv6 PIM Sparse routers within the IPv6 PIM Sparse domain. Make sure the device is on the backbone or is otherwise well-connected to the rest of the network.

To specify the IPv6 address of the RP, enter commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# rp-address 31::207
```

The command in the previous example identifies the router interface at IPv6 address 31:207 as the RP for the IPv6 PIM Sparse domain. The device will use the specified RP and ignore group-to-RP mappings received from the BSR.

To specify the IPv6 address of the RP for a specified VRF, enter commands such as the following.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# rp-address 31::207
```

Syntax: `[no] rp-address ipv6-addr`

The `ipv6-addr` parameter specifies the IPv6 address of the RP.

Updating IPv6 PIM Sparse forwarding entries with a new RP configuration

If you make changes to your static RP configuration, the entries in the IPv6 PIM Sparse multicast forwarding table continue to use the old RP configuration until they are aged out.

The `clear IPv6 pim rp-map` command allows you to update the entries in the static multicast forwarding table immediately after making RP configuration changes. This command is meant to be used with the `rp-address` command.

To update the entries in an IPv6 PIM Sparse static multicast forwarding table with a new RP configuration, enter the following command at the privileged EXEC level of the CLI.

```
device(config)# clear ipv6 pim rp-map
```

Syntax: `clear ipv6 pim rp-map`

Embedded Rendezvous Point

Global deployment of IPv4 multicast relies on Multicast Source Discovery Protocol (MSDP) to convey information about the active sources. Because IPv6 provides more address space, the RP address can be included in the multicast group address.

NOTE

The IPv6 group address must be part of the FF70:/12 prefix.

Embedded RP support is enabled by default. You can disable it using the following commands.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# no rp-embedded
```

To disable embedded RP support for a specified VRF, enter the following commands.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# no rp-embedded
```

Syntax: [no] rp-embedded

Changing the Shortest Path Tree threshold

In a typical IPv6 PIM Sparse domain, there may be two or more paths from a designated router (DR) for a multicast source to an IPv6 PIM group receiver:

- **Path through the RP** - This is the path the device uses the first time it receives traffic for an IPv6 PIM group. However, the path through the RP may not be the shortest path from the device to the receiver.
- **Shortest Path** - Each IPv6 PIM Sparse router that is a DR for an IPv6 receiver calculates a short path tree (SPT) towards the source of the IPv6 multicast traffic. The first time the device configured as an IPv6 PIM router receives a packet for an IPv6 group, it sends the packet to the RP for that group, which in turn will forward it to all the intended DRs that have registered with the RP. The first time the device is a recipient, it receives a packet for an IPv6 group and evaluates the shortest path to the source and initiates a switchover to the SPT. Once the device starts receiving data on the SPT, the device proceeds to prune itself from the **RPT**.

By default, the device switches from the RP to the SPT after receiving the first packet for a given IPv6 PIM Sparse group. The device maintains a separate counter for each IPv6 PIM Sparse source-group pair.

You can change the number of packets the device receives using the RP before switching to using the SPT.

To change the number of packets the device receives using the RP before switching to the SPT, enter commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# spt-threshold 1000
```

To change the number of packets the device receives using the RP before switching to the SPT for a specified VRF, enter commands such as the following.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# spt-threshold 1000
```

Syntax: [no] spt-threshold infinity | num

The **infinity** | num parameter specifies the number of packets. If you specify **infinity**, the device sends packets using the RP indefinitely and does not switch over to the SPT. If you enter a specific number of packets, the device does not switch over to using the SPT until it has sent the number of packets you specify using the RP.

Setting the RP advertisement interval

To specify how frequently the candidate RP configured on the device sends candidate RP advertisement messages to the BSR, enter commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# rp-adv-interval 180
Changed RP ADV interval to 180 seconds.
```

To specify how frequently the candidate RP configured on the device sends candidate RP advertisement messages to the BSR for a specified VRF, enter commands such as the following.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# rp-adv-interval 180
Changed RP ADV interval to 180 seconds.
```

Syntax: `rp-adv-interval seconds`

The *seconds* parameter specifies the number of seconds. The default is 60 seconds.

Route selection precedence for multicast

The **route-precedence** command allows the user to specify a precedence table that dictates how routes are selected for multicast.

NOTE

PIM must be enabled at the global level.

Configuring the route precedence by specifying the route types

The **route-precedence** *mc-non-default mc-default uc-non-default uc-default none* command allows you to control the selection of routes based on the route types. There are four different types of routes:

- Non-default route from the mRTM
- Default route from the mRTM
- Non-default route from the uRTM
- Default route from the uRTM

Using the **route-precedence** command, you may specify an option for all of the precedence levels.

To specify a non-default route from the mRTM, then a non-default route from the uRTM, then a default route from the mRTM, and then a default route from the uRTM, enter commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# route-precedence mc-non-default uc-non-default mc-default uc-default
```

The *none* option can be used to fill up the precedence table in order to ignore certain types of routes. To use the unicast default route for multicast, enter commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# route-precedence mc-non-default mc-default uc-non-default none
```

To specify a non-default route from the mRTM, then a non-default route from the uRTM, then a default route from the mRTM, and then a default route from the uRTM for a specified VRF, enter commands such as the following.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# route-precedence mc-non-default uc-non-default mc-default uc-default
```

The *none* option can be used to fill up the precedence table in order to ignore certain types of routes. To use the unicast default route for multicast for a specified VRF, enter commands such as the following.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# route-precedence mc-non-default mc-default uc-non-default none
```

Syntax: **[no] route-precedence { mc-non-default | mc-default | uc-non-default | uc-default | none }**

The default value is the **route-precedence mc-non-default mc-default uc-non-default uc-default** command.

Use the **mc-non-default** parameter to specify a multicast non-default route.

Use the **mc-default** parameter to specify a multicast default route.

Use the **uc-non-default** parameter to specify a unicast non-default route.

Use the **uc-default** parameter to specify a unicast default route.

Use the **none** parameter to ignore certain types of routes.

The **no** form of this command removes the configuration.

Changing the PIM Join and Prune message interval

By default, the device sends PIM Sparse Join or Prune messages every 60 seconds. These messages inform other PIM Sparse routers about clients who want to become receivers (Join) or stop being receivers (Prune) for PIM Sparse groups.

NOTE

Use the same Join or Prune message interval on all the PIM Sparse routers in the PIM Sparse domain. If the routers do not all use the same timer interval, the performance of PIM Sparse can be adversely affected.

To change the Join or Prune interval, enter commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# message-interval 30
```

To change the Join or Prune interval for a specified VRF, enter the commands as shown in the following example.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# message-interval 30
```

Syntax: **[no] message-interval seconds**

The *seconds* parameter specifies the number of seconds and can be from 1 through 65535 seconds. The default is 60 seconds.

Modifying neighbor timeout

Neighbor timeout is the interval after which a PIM router will consider a neighbor to be absent. If the timer expires before receiving a new hello message, the PIM router will time out the neighbor.

To apply an IPv6 PIM neighbor timeout value of 33 seconds to all ports on the router operating with PIM, enter the commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# nbr-timeout 33
```

To apply an IPv6 PIM neighbor timeout value of 33 seconds for a specified VRF operating with PIM, enter the commands such as the following.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# nbr-timeout 33
```

Syntax: `[no] nbr-timeout seconds`

The *seconds* parameter specifies the number of seconds. The valid range is from 35 through 65535 seconds.

Setting the prune wait interval

The **prune-wait** command allows you to set the amount of time the PIM router should wait for a join override before pruning an Outgoing Interface List Optimization (OIF) from the entry.

To change the default join override time to 2 seconds, enter commands such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# prune-wait 2
```

To change the default join override time to 2 seconds for a specified VRF, enter commands such as the following.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# prune-wait 2
```

Syntax: `[no] prune-wait seconds`

The *seconds* parameter specifies the number of seconds. The valid range is from 0 through 30 seconds. The default is 3 seconds.

Setting the register suppress interval

The **register-suppress-time** command allows you to set the amount of time the PIM router uses to periodically trigger the NULL register message.

NOTE

The register suppress time configuration applies only to the first hop PIM router.

To change the default register suppress time to 90 seconds, enter commands such as the following:

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# register-suppress-time 90
```

To change the default register suppress time to 90 seconds for a specified VRF, enter commands such as the following:

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# register-suppress-time 90
```

Syntax: `[no] register-suppress-time seconds`

The *seconds* parameter specifies the number of seconds. The valid range is from 60 through 120 seconds. The default is 60 seconds.

Setting the register probe time

The **register-probe-time** command allows you to set the amount of time the PIM router waits for a register-stop from an RP before it generates another NULL register to the PIM RP. The register probe time configuration applies only to the first hop PIM router.

NOTE

Once a PIM first hop router successfully registers with a PIM RP, the PIM first hop router will not default back to the data registration. All subsequent registers will be in the form of the NULL registration.

To change the default register probe time to 20 seconds, enter commands such as following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# register-probe-time 20
```

To change the default register probe time to 20 seconds for a specified VRF, enter commands such as the following.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# register-probe-time 20
```

Syntax: [no] **register-probe-time** *seconds*

The *seconds* parameter specifies the number of seconds. The valid range is from 10 through 50 seconds. The default is 10 seconds.

Setting the inactivity timer

The router deletes a forwarding entry if the entry is not used to send multicast packets. The IPv6 PIM inactivity timer defines how long a forwarding entry can remain unused before the router deletes it.

To apply an IPv6 PIM inactivity timer of 160 seconds to all IPv6 PIM interfaces, enter the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# inactivity-timer 160
```

To apply an IPv6 PIM inactivity timer of 160 seconds for a specified VRF, enter the commands as shown in the following example.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# inactivity-timer 160
```

Syntax: [no] **inactivity-timer** *seconds*

The *seconds* parameter specifies the number of seconds. The valid range is 60 through 3600 seconds. The default is 180 seconds.

Changing the hello timer

The hello timer defines the interval at which periodic hellos are sent out to PIM interfaces. Routers use hello messages to inform neighboring routers of their presence. To change the hello timer, enter a command such as the following.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# hello-timer 62
```

To change the hello timer for a specified VRF, enter the commands as shown in the following example.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# hello-timer 62
```

Syntax: [no] **hello-timer** *seconds*

The *seconds* parameter specifies the number of seconds. The valid range is 10 through 3600 seconds. The default is 60 seconds.

Enabling Source-specific Multicast

Using the Any-Source Multicast (ASM) service model, sources and receivers register with a multicast address. The protocol uses regular messages to maintain a correctly configured broadcast network where all sources can send data to all receivers and all receivers get broadcasts from all sources.

With Source-specific Multicast (SSM), the "channel" concept is introduced where a "channel" consists of a single source and multiple receivers that specifically register to get broadcasts from that source. Consequently, receivers are not burdened with receiving data they have no interest in, and network bandwidth requirements are reduced because the broadcast need only go to a subset of users. The address range ff30:/12 has been assigned by the Internet Assigned Numbers Authority (IANA) for use with SSM.

SSM simplifies IPv6 PIM-SM by eliminating the RP and all protocols related to the RP.

Configuring Source-specific Multicast

IPv6 PIM-SM must be enabled on any ports on which you want SSM to operate. Enter the **ssm-enable** command under the IPv6 router PIM level to globally enable SSM filtering.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# ssm-enable ff02::200:2
```

To enable SSM for a specified VRF, enter the commands as shown in the following.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# ssm-enable ff02::200:2
```

Syntax: **[no] ssm-enable [range address-range]**

The **rangeaddress-range** option allows you to define the SSM range of IPv6 multicast addresses.

Modifying the Hop-Limit threshold

The Time To Live (TTL) defines the minimum value required in a packet in order for the packet to be forwarded out the interface. For example, if the Hop-Limit for an interface is set at 10, it means that only those packets that ingress with a TTL value of 11 or more will be forwarded out the TTL-10 interface. Thus, with a default Hop-Limit threshold of 1, only packets ingressing with a Hop-Limit of 2 or greater will be forwarded out. Note that the Hop-Limit threshold only applies to routed interfaces. Switched interfaces ignore the Hop-Limit threshold. Possible Hop-Limit values are from 1 through 64. The default Hop-Limit value is 1. To configure a Hop-Limit of 45, enter the following command.

```
device(config-if-e10000-3/24)# ipv6 pim ttl-threshold 45
```

To configure a Hop-Limit of 45 on a virtual Ethernet interface, enter the following commands.

```
device(config)# interface ve 10
device(config-vif-10)# ipv6 pim ttl-threshold 45
```

Syntax: **ipv6 pim ttl-threshold 1-64**

Configuring a DR priority

The DR priority option lets a network administrator give preference to a particular router in the DR election process by giving it a numerically higher DR priority. To set a DR priority higher than the default value of 1, use the **ipv6 pim dr-priority** command as shown in the example below.

```
device(config-if-e10000-3/24)# ipv6 pim dr-priority 50
```

To set a DR priority higher than the default value of 1 on a virtual Ethernet interface, use the **ipv6 pim dr-priority** command as shown in the following.

```
device(config)# interface ve 10
device(config-vif-10)# ipv6 pim dr-priority 50
```

Syntax: **[no] ipv6 pim dr-priority priority-value**

The *priority-value* variable is the value that you want to set for the DR priority. The range of values is from 0 through 65535. The default value is 1.

The **no** option removes the command and sets the DR priority back to the default value of 1.

The following information may be useful for troubleshooting:

- If more than one router has the same DR priority on a subnet (as in the case of default DR priority on all), the router with the numerically highest IP address on that subnet will get elected as the DR.

- The DR priority information is used in the DR election *only if all* the PIM routers connected to the subnet support the DR priority option. If there is at least one PIM router on the subnet that does not support this option, then the DR election falls back to the backwards compatibility mode in which the router with the numerically highest IP address on the subnet is declared the DR regardless of the DR priority values.

Passive Multicast Route Insertion

To prevent unwanted multicast traffic from being sent to the CPU, IPv6 PIM routing and Passive Multicast Route Insertion (PMRI) can be used together to ensure that multicast streams are only forwarded out ports with interested receivers and unwanted traffic is dropped in hardware on Layer 3 routers.

PMRI enables a Layer 3 switch running IPv6 PIM Sparse to create an entry for a multicast route (for example, (S,G)), with no directly attached clients or when connected to another PIM router (transit network).

When a multicast stream has no output interfaces, the Layer 3 switch can drop packets in hardware if the multicast traffic meets the following conditions in IPv6 PIM-SM.

- The route has no OIF.
- The directly connected source passes source RPF check and completes data registration with the RP, or the non-directly connected source passes source RPF check.

If the OIF is inserted after the hardware-drop entries are installed, the hardware entries will be updated to include the OIFs.

NOTE

Disabling hardware-drop does not immediately take away existing hardware-drop entries, they will go through the normal route aging processing when the traffic stops.

Configuring PMRI

PMRI is enabled by default. To disable PMRI, enter the following commands.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# hardware-drop-disable
```

To disable PMRI for a specified VRF, enter the commands as shown in the following example.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# hardware-drop-disable
```

Syntax: [no] hardware-drop-disable

Displaying hardware-drop

Use the **show ipv6 pim sparse** command to display if the hardware-drop feature has been enabled or disabled.

```
device# show ipv6 pim sparse
Global PIM Sparse Mode Settings
  Hello interval          : 30           Neighbor timeout          : 105
  Bootstrap Msg interval: 60           Candidate-RP Advertisement interval: 60
  Join/Prune interval    : 60           SPT Threshold            : 1
  SSM Enabled: Yes
  SSM Group Range: ff30::/12
  Hardware Drop Enabled : Yes
```

Displaying PIM Sparse configuration information and statistics

You can display the following PIM Sparse information:

- Basic PIM Sparse configuration information
- IPv6 interface information
- Group information
- BSR information
- Candidate RP information
- RP-to-group mappings
- RP information for an IPv6 PIM Sparse group
- RP set list
- Multicast neighbor information
- The IPv6 PIM multicast cache
- IPv6 PIM RPF
- IPv6 PIM counters
- IPv6 PIM resources
- IPv6 PIM traffic statistics

Displaying basic PIM Sparse configuration information

To display IPv6 PIM Sparse configuration information, enter the **show ipv6 pim sparse** command at any CLI level.

```
device show ipv6 pim sparse
Global PIM Sparse Mode Settings
  Hello interval           : 30           Neighbor timeout           : 105
  Bootstrap Msg interval  : 60           Candidate-RP Advertisement interval: 60
  Register Suppress interval: 60         Register Stop Delay       : 60
  Join/Prune interval     : 60           SPT Threshold             : 1
  Inactivity interval     : 180          Hardware Drop Enabled     : Yes
  SSM Enabled              : Yes
```

Syntax: `show ipv6 pim [vrf vrf-name] sparse`

The *vrf* parameter allows you to configure IPv6 PIM on the virtual routing instance (VRF) specified by the *vrf-name* variable.

[Table 31](#) displays the output from the **show ipv6 pim sparse** command.

TABLE 31 Output from the **show ipv6 pim sparse** command

Field	Description
Global PIM Sparse mode settings	
Hello interval	How frequently the device sends IPv6 PIM Sparse hello messages to its IPv6 PIM Sparse neighbors. This field shows the number of seconds between hello messages. IPv6 PIM Sparse routers use hello messages to discover one another.
Neighbor timeout	How many seconds the device will wait for a hello message from a neighbor before determining that the neighbor is no longer present and removing cached IPv6 PIM Sparse forwarding entries for the neighbor.
Bootstrap Msg interval	How frequently the BSR configured on the device sends the RP set to the RPs within the IPv6 PIM Sparse domain. The RP set is a list of candidate RPs and their group prefixes. The group prefix of a candidate RP indicates the range of IPv6 PIM Sparse group numbers for which it can be an RP.

TABLE 31 Output from the `show ipv6 pim sparse` command (continued)

Field	Description
	<p>NOTE This field contains a value only if an interface on the device is elected to be the BSR. Otherwise, the field is blank.</p>
Candidate-RP Advertisement interval	<p>How frequently the candidate RP configured on the device sends candidate RP advertisement messages to the BSR.</p> <p>NOTE This field contains a value only if an interface on the device is configured as a candidate RP. Otherwise, the field is blank.</p>
Join or Prune interval	<p>How frequently the device sends IPv6 PIM Sparse Join or Prune messages for the multicast groups it is forwarding. This field shows the number of seconds between Join or Prune messages.</p> <p>The device sends Join or Prune messages on behalf of multicast receivers that want to join or leave an IPv6 PIM Sparse group. When forwarding packets from IPv6 PIM Sparse sources, the device sends the packets only on the interfaces on which it has received join requests in Join or Prune messages for the source group.</p>
SPT Threshold	The number of packets the device sends using the path through the RP before switching to using the SPT path.
Inactivity Interval	How long a forwarding entry can remain unused before the router deletes it.
SSM Enabled	If yes, source-specific multicast is configured globally on this router.
IPv6 PIM Sparse interface information	
	<p>NOTE You also can display IPv6 multicast interface information using the <code>show ipv6 pim interface</code> command.</p>
Interface	<p>The type of interface and the interface number. The interface type can be one of the following:</p> <ul style="list-style-type: none"> Ethernet VE <p>The number is either a port number (and slot number if applicable) or the virtual interface (VE) number.</p>
TTL Threshold	<p>Following the TTL threshold value, the interface state is listed. The interface state can be one of the following:</p> <ul style="list-style-type: none"> Disabled Enabled
Local Address	Indicates the IP address configured on the port or virtual interface.

Displaying IPv6 PIM interface information

You can display IPv6 PIM multicast interface information using the `show ipv6 pim interface` command.

To display IPv6 PIM multicast interface information, enter the following command as shown in the example.

```

device show ipv6 pim interface
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Inter|Global Address      |Ver|St |TTL|Multicast|Filter|VRF| DR |Override
face|+ Designated Router Port|  |  |Thr|filter  |ACL  |  |Prio|Interval
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
v55  2001:10::1          SMv2 Ena 1  None  None  Default 1  3000ms
+ Itself

```

```
v56 2001:2::1 SMv2 Ena 1 None None Default 1 3000ms
+ fe80::224:38ff:fe9b:ce00
3/1
```

Syntax: `show ipv6 pim [vrf vrf-name] interface [ethernet slot/port | loopback number | ve number]`

The **vrf** parameter allows you to display IPv6 multicast interface information for the VRF instance identified by the *vrf-name* variable.

The **ethernet slot/portnum | loopback num | ve num** parameter specifies the IPv6 PIM multicast interface.

- Enter **ethernet slot/portnum** for a physical interface (port).
- Enter **loopback num** for a loopback interface.
- Enter **ve num** for a virtual interface.

Displaying a list of multicast groups

To display IPv6 PIM group information, enter the **show ipv6 pim group** command at any CLI level.

```
device show ipv6 pim group
Total number of groups: 1
1 Group ff7e:a40:2001:3e8:27:0:1:2 Ports
Group member at e3/1: v31
```

Syntax: `show ipv6 pim [vrf vrf-name] group`

The **vrf** parameter allows you to display IPv6 PIM group information for the VRF instance identified by the *vrf-name* variable.

[Table 32](#) displays the output from the **show ipv6 pim group** command.

TABLE 32 Output from the **show ipv6 pim group** command

Field	Description
Total number of Groups	Lists the total number of IPv6 multicast groups the device is forwarding.
Group	The multicast group address.
Ports	The device ports connected to the receivers of the groups.

Displaying BSR information

To display information on a device that has been elected as the BSR, enter the **show ipv6 pim bsr** command at the CLI level.

```
device show ipv6 pim bsr
PIMv2 Bootstrap information

This system is the elected Bootstrap Router (BSR)
BSR address: 2001:3e8:255:255::17
Uptime: 00:12:09, BSR priority: 0, Hash mask length: 126
Next bootstrap message in 00:00:30

Next Candidate-RP-advertisement in 00:00:30
RP: 2001:3e8:255:255::17
group prefixes:
ff00:: / 8

Candidate-RP-advertisement period: 60
```

The following example shows information displayed on a device that is not the BSR. Notice that some fields shown in the previous example do not appear in the following example.

```
device# show ipv6 pim bsr
PIMv2 Bootstrap information
BSR address = 2001:3e8:255:255::17
BSR priority = 0
```

Syntax: `show ipv6 pim [vrf vrf-name] bsr`

The `vrf` parameter allows you to display IPv6 PIM BSR information for the VRF instance identified by the `vrf-name` variable.

Table 33 displays the output from the `show ipv6 pim bsr` command.

TABLE 33 Output from the `show ipv6 pim bsr` command

Field	Description
BSR address	The IPv6 address of the interface configured as the IPv6 PIM Sparse Bootstrap Router (BSR).
Uptime	The amount of time the BSR has been running. NOTE This field appears only if this device is the BSR.
BSR priority	The priority assigned to the interface for use during the BSR election process. During BSR election, the priorities of the candidate BSRs are compared and the interface with the highest BSR priority becomes the BSR.
Hash mask length	The number of significant bits in the IPv6 multicast group comparison mask. This mask determines the IPv6 multicast group numbers for which the device can be a BSR. The default is 32 bits, which allows the device to be a BSR for any valid IPv6 multicast group number. NOTE This field appears only if this device is a candidate BSR.
Next bootstrap message in	Indicates how many seconds will pass before the BSR sends its next Bootstrap message. NOTE This field appears only if this device is the BSR.
Next Candidate-RP-advertisement message in	Indicates how many seconds will pass before the BSR sends its next candidate RP advertisement message. NOTE This field appears only if this device is a candidate BSR.
RP	Indicates the IPv6 address of the Rendezvous Point (RP). NOTE This field appears only if this device is a candidate BSR.
group prefixes	Indicates the multicast groups for which the RP listed by the previous field is a candidate RP. NOTE This field appears only if this device is a candidate BSR.
Candidate-RP-advertisement period	Indicates how frequently the BSR sends candidate RP advertisement messages.

TABLE 33 Output from the `show ipv6 pim bsr` command (continued)

Field	Description
	<p>NOTE This field appears only if this device is a candidate BSR.</p>

Displaying candidate RP information

To display candidate RP information, enter the `show ipv6 rp-candidate` command at any CLI level.

```
device# show ipv6 pim rp-candidate
Next Candidate-RP-advertisement in 00:00:10
  RP: 1be::11:21
    group prefixes:
      ff00:: / 8
Candidate-RP-advertisement period: 60
```

This example shows information displayed on a device that is a candidate RP. The following example shows the message displayed on a device that is not a candidate RP.

```
device# show ipv6 pim rp-candidate
```

This system is not a Candidate-RP.

Syntax: `show ipv6 pim [vrf vrf-name] rp-candidate`

The `vrf` parameter allows you to display IPv6 candidate RP information for the VRF instance identified by the `vrf-name` variable.

Table 34 displays the output from the `show ipv6 pim rp-candidate` command.

TABLE 34 Output from the `show ipv6 pim rp-candidate` command

Field	Description
Candidate-RP-advertisement in	<p>Indicates how many seconds will pass before the BSR sends its next RP message.</p> <p>NOTE This field appears only if this device is a candidate RP.</p>
RP	<p>Indicates the IPv6 address of the Rendezvous Point (RP).</p> <p>NOTE This field appears only if this device is a candidate RP.</p>
group prefixes	<p>Indicates the multicast groups for which the RP listed by the previous field is a candidate RP.</p> <p>NOTE This field appears only if this device is a candidate RP.</p>
Candidate-RP-advertisement period	<p>Indicates how frequently the BSR sends candidate RP advertisement messages.</p> <p>NOTE This field appears only if this device is a candidate RP.</p>

Displaying RP-to-group mappings

To display RP-to-group-mappings, enter the **show ipv6 pim rp-map** command at any CLI level.

```
device# show ipv6 pim rp-map
Idx Group address                      RP address
 1 fff1e::1:2 2001:3e8:255:255::17
 2                ff7e:a40:2001:3e8:27:0:1:2 2001:3e8:27::a
 3                ff7e:140:2001:3e8:16:0:1:2 2001:3e8:16::1
```

Syntax: **show ipv6 pim [vrf vrf-name] rp-map**

The **vrf** parameter allows you to display IPv6 RP-to-group-mappings for the VRF instance identified by the *vrf-name* variable.

[Table 35](#) displays the output from the **show ipv6 rp-map** command.

TABLE 35 Output from the **show ipv6 pim rp-map** command

Field	Description
Index	The index number of the table entry in the display.
Group address	Indicates the IPv6 PIM Sparse multicast group address using the listed RP.
RP address	Indicates the IPv6 address of the Rendezvous Point (RP) for the listed PIM Sparse group.

Displaying RP information for an IPv6 PIM Sparse group

To display RP information for an IPv6 PIM Sparse group, enter the following command at any CLI level.

```
device# show ipv6 pim rp-hash ffile::1:2
RP: 2001:3e8:255:255::17, v2
Info source: 2001:3e8:255:255::17, via bootstrap
```

Syntax: **show ipv6 pim [vrf vrf-name] rp-hash group-addr**

The **vrf** parameter allows you to display RP information for a PIM Sparse group for the VRF instance identified by the *vrf-name* variable.

The *group-addr* parameter is the address of an IPv6 PIM Sparse IP multicast group.

[Table 36](#) displays the output from the **show ipv6 pim rp-hash group-addr** command.

TABLE 36 Output from the **show ipv6 pim rp-hash group-addr** command

Field	Description
RP	Indicates the IPv6 address of the Rendezvous Point (RP) for the specified IPv6 PIM Sparse group. Following the IPv6 address is the port or virtual interface through which this device learned the identity of the RP.
Info source	Indicates the IPv6 address on which the RP information was received. Following the IPv6 address is the method through which this device learned the identity of the RP.

Displaying the RP set list

To display the RP set list, enter the **show ipv6 pim rp-set** command at any CLI level.

```
device# show ipv6 pim rp-set
Static RP
-----
Static RP count: 1
100::1
```

```
Number of group prefixes Learnt from BSR: 0
No RP-Set present
```

Syntax: `show ipv6 pim [vrf vrf-name] rp-set`

The `vrf` parameter allows you to display the RP set for the VRF instance identified by the `vrf-name` variable.

[Table 37](#) displays the output from the `show ipv6 pim rp-set` command.

TABLE 37 Output from the `show ipv6 pim rp-set` command

Field	Description
Number of group prefixes	The number of IPv6 PIM Sparse group prefixes for which the RP is responsible.
Group prefix	Indicates the multicast groups for which the RP listed by the previous field is a candidate RP.
RPs expected or received	Indicates how many RPs were expected and received in the latest Bootstrap message.
RP num	Indicates the RP number. If there are multiple RPs in the IPv6 PIM Sparse domain, a line of information for each of them is listed, and they are numbered in ascending numerical order.
priority	The RP priority of the candidate RP. During the election process, the candidate RP with the highest priority is elected as the RP.
age	The age (in seconds) of this RP-set. NOTE If this device is not a BSR, this field contains zero. Only the BSR ages the RP-set.

Displaying multicast neighbor information

To display information about IPv6 PIM neighbors, enter the `show ipv6 pim neighbor` command at any CLI level.

```
device# show ipv6 pim neighbor
Port Phy_Port Neighbor                Holdtime Age    UpTime Priority
      sec      sec      sec
e11/15 e11/15 fe80::45:27:49:4 105      20    1010     1
v312   e11/3   fe80::45:27:1:2   105      10    1900    40
```

Syntax: `show ipv6 pim [vrf vrf-name] neighbor`

The `vrf` parameter allows you to display the IPv6 PIM neighbors for the VRF instance identified by the `vrf-name` variable.

[Table 38](#) displays the output from the `show ipv6 pim neighbor` command.

TABLE 38 Output from the `show ipv6 pim neighbor` command

Field	Description
Port	The interface through which the device is connected to the neighbor.
Neighbor	The IPv6 interface of the IPv6 PIM neighbor interface.
Holdtime sec	Indicates how many seconds the neighbor wants this device to hold the entry for this neighbor in memory. The neighbor sends the Hold Time in its hello packets. <ul style="list-style-type: none"> If the device receives a new hello packet before the Hold Time received in the previous packet expires, the device updates its table entry for the neighbor. If the device does not receive a new hello packet from the neighbor before the Hold time expires, the device assumes the

TABLE 38 Output from the `show ipv6 pim neighbor` command (continued)

Field	Description
	neighbor is no longer available and removes the entry for the neighbor.
Age sec	The number of seconds since the device received the last hello message from the neighbor.
UpTime sec	The number of seconds the PIM neighbor has been up. This timer starts when the device receives the first hello messages from the neighbor.
Priority	The DR priority that is used in the DR election process. This can be a configured value or the default value of 1.

Displaying the IPv6 PIM multicast cache

To display the IPv6 PIM multicast cache, enter the `show ipv6 pim mcache` command at any CLI level.

NOTE

Brocade NetTron CES and NetTron CER devices display incorrect hardware programmed entries. The information displayed for the forwarding port should be disregarded.

```

device#show ipv6 pim vrf sd2 mcache
Total entries in mcache: 4
 1  (*, ffle::1) RP 2005::192:168:1:1, in NIL (NIL), Uptime 00:27:05
    Sparse Mode, RPT=1 SPT=0 Reg=0 L2Reg=0 RegSupp=0 RegProbe=0 LSrc=0 LRcv=1
    No upstream neighbor because RP 2005::192:168:1:1 is itself
    num_oifs = 1
    immediate_oifs = 1, inherited_oifs = 0, blocked_oifs = 0
      152,4/12(00:27:05/0) Flags:00000004
    slow ports ethe 4/12
    L3 (SW) 1: e4/12(VL152)
    Flags (0x00260480)
      sm=1 ssm=0 hw=0 fast=0 slow=0 leaf=0 prun=0 tag=0 needRte=0 msdp_adv=0
    AgeSltMsk=00000000, FID: NotReq MVID: NotReq, RegPkt 0 profile: 2
 2  (2001:1:1:1::149:101, ffle::1) in v149 (tag e4/9), Uptime 00:27:43 Rate 2284
    Sparse Mode, RPT=0 SPT=1 Reg=0 L2Reg=1 RegSupp=0 RegProbe=0 LSrc=1 LRcv=1
    Source is directly connected. RP 2005::192:168:1:1
    num_oifs = 1
    immediate_oifs = 0, inherited_oifs = 1, blocked_oifs = 0
      152,4/12(00:27:05/0) Flags:00000004
    fast ports ethe 4/12
    L3 (HW) 1: e4/12(VL152)
    Flags (0x604688e1)
      sm=1 ssm=0 hw=1 fast=1 slow=0 leaf=0 prun=0 tag=1 needRte=0 msdp_adv=0
    AgeSltMsk=00000008, FID: 0x800f MVID: 1, RegPkt 0, AvgRate 2228 profile: 2
 3  (*, ffle::2) RP 2005::192:168:1:1, in NIL (NIL), Uptime 00:27:05
    Sparse Mode, RPT=1 SPT=0 Reg=0 L2Reg=0 RegSupp=0 RegProbe=0 LSrc=0 LRcv=1
    No upstream neighbor because RP 2005::192:168:1:1 is itself
    num_oifs = 1
    immediate_oifs = 1, inherited_oifs = 0, blocked_oifs = 0
      152,4/12(00:27:05/0) Flags:00000004
    slow ports ethe 4/12
    L3 (SW) 1: e4/12(VL152)
    Flags (0x00260480)
      sm=1 ssm=0 hw=0 fast=0 slow=0 leaf=0 prun=0 tag=0 needRte=0 msdp_adv=0
    AgeSltMsk=00000000, FID: NotReq MVID: NotReq, RegPkt 0 profile: 2
 4  (2001:1:1:1::149:101, ffle::2) in v149 (tag e4/9), Uptime 00:27:43 Rate 2284
    Sparse Mode, RPT=0 SPT=1 Reg=0 L2Reg=1 RegSupp=0 RegProbe=0 LSrc=1 LRcv=1
    Source is directly connected. RP 2005::192:168:1:1
    num_oifs = 1
    immediate_oifs = 0, inherited_oifs = 1, blocked_oifs = 0
      152,4/12(00:27:05/0) Flags:00000004
    fast ports ethe 4/12
    L3 (HW) 1: e4/12(VL152)
    Flags (0x604688e1)

```

```
sm=1 ssm=0 hw=1 fast=1 slow=0 leaf=0 prun=0 tag=1 needRte=0 msdp_adv=0
AgeSltMsk=00000008, FID: 0x8010 MVID: 1, RegPkt 0, AvgRate 2228 profile: 2
```

Syntax: `show ipv6 pim mcache [multicast cache entries source/groupaddress | multicast cacheipv6-group-address]`

Syntax: `show ipv6 pim [vrf vrf-name] mcache`

The `vrf` parameter allows you to display the IPv6 PIM multicast cache for the VRF instance identified by the `vrf-name` variable.

[Table 39](#) describes the output parameters of the `show ipv6 pim vrf mcache` command.

TABLE 39 Output parameters of the `show ipv6 pim vrf mcache` command

Field	Description
Total entries in mcache	Shows the total number of PIM mcache entries.
Uptime	Shows the entry uptime.
Rate	Shows the total packet count.
Sparse mode	Shows that the PIM sparse mode is enabled.
RPT	Sets the flag to 1, if the entry is on the RP tree, else sets the flag to 0.
SPT	Sets the flag to 1, if the entry is on the source tree, or else sets the flag to 0.
Reg	Sets the flag to 1, if the data registration is in progress, or else sets the flag to 0.
L2Reg	Sets the flag to 1, if the source is directly connected to the router, or else sets the flag to 0.
RegSupp	Sets the flag to 1, if the register suppression timer is running, or else sets the flag to 0.
RegProbe	Sets the flag to 1, if the mcache entry is entering the register probing period, or else sets the flag to 0.
LSrc	Sets the flag to 1, if the source is in a directly-connected interface, or else sets the flag to 0.
LRcv	Sets the flag to 1, if the receiver is directly connected to the router, or else sets the flag to 0.
RP	Show the IP address of the RP.
num_oifs	Show the count of the outgoing interfaces.
inherited_oifs	Shows the PIM Sparse mode inherited outgoing interfaces.
blocked_oifs	Show the PIM Sparse mode blocked outgoing interfaces.
immediate_oifs	Shows the local immediate outgoing interface of the mcache entry.
Flags	Show the flags associated with the forward entry.
slow ports ethe	Shows the forwarding port ID of the mcache entry which is in the software forwarding path.
L3	Show whether the traffic is switched or routed out of the interface.
(HW)	Show whether the entry is software forwarded or hardware forwarded.
sm	Sets the flag to 1, if the Sparse Mode entry is enabled and 0, if the PIM dense mode entry is enabled.
ssm	Sets the flag to 1, if the SSM mode entry is enabled, or else sets the flag to 0.
hw	Sets the flag to 1, if the candidate for hardware forwarding is enabled, or else sets the flag to 0.
fast	Sets the flag to 1, if the resources are allocated for hardware forwarding, or else sets the flag to 0.

TABLE 39 Output parameters of the `show ipv6 pim vrf mcache` command (continued)

Field	Description
slow	Sets the flag to 1, if the entry is not a candidate for hardware forwarding or the resource allocation is failed, or else sets the flag to 0.
prun	Sets the flag to 1, if PIM is enabled, or else sets the flag to 0.
tag	Sets the flag to 1, if there is a need for allocating entries from the replication table, or else sets the flag to 0.
needRte	Sets the flag to 1, if there is no route to the source and RP is available, or else sets the flag to 0.
msdp_adv	Sets the flag to 1, if RP is responsible for the source and must be advertised to its peers.
AgeSlitMsk	Shows the slot number on which MP expects ingress traffic.
FID	Shows the FID resource allocated for a particular entry.
MVID	Shows the MVID resource allocated for a particular entry.
RegPkt	Shows the number of PIM register packet received.
AvgRate	Shows the average data traffic rate for the mcache entry
profile	Shows the profile ID associated with the stream.

Displaying IPv6 PIM RPF

The `show ipv6 pim rpf` command displays what PIM sees as the reverse path to the source. While there may be multiple routes back to the source, the one displayed by the `show ipv6 pim rpf` command is the one that PIM thinks is best.

```
device# show ipv6 pim vrf eng rpf 130.50.11.10
Source 130.50.11.10 directly connected on e4/1
```

Syntax: `show ipv6 pim [vrf vrf-name] rpf ip-address`

The `vrf` parameter allows you to display what PIM sees as the reverse path to the source for a VRF instance specified by the `vrf-name` variable.

The `ip-address` variable specifies the source address for RPF check.

Displaying IPv6 PIM counters

You can display the number of default-vlan-id changes that have occurred since the applicable VRF was created and how many times a tagged port was placed in a VLAN since the applicable VRF was created, as shown in the following example.

```
device(config)# show ipv6 pim vrf eng counter
Event Callback:
  DFTVlanChange      :          0          VlanPort      :          2
LP to MP IPCs:
  SM_REGISTER        :          0          MCAST_CREATE    :          31
  S_G_AGEOUT         :          4          WRONG_IF        :          0
  ABOVE_THRESHOLD    :          0          MCAST_FIRST_DATA :          31
```

Syntax: `show ipv6 pim [vrf vrf-name] counter`

The `vrf` parameter allows you to display IPv6 PIM counters for the VRF instance identified by the `vrf-name` variable.

Table 40 displays the output from the `show ipv6 vrf eng counter` command.

TABLE 40 Output from the `show ipv6 pim vrf eng counter` command

Field	Description
DFTVlanChange	The number of default-vlan-id changes that have occurred since the applicable VRF was created.
VlanPort	The number of times that a tagged port was placed in a VLAN since the applicable VRF was created.

Displaying the IPv6 PIM resources

To display the hardware resource information, such as hardware allocation, availability, and limit for software data structure, enter the `show ipv6 pim resource` command.

```

device# show ipv6 pim resource
              allocated   in-use  available  allo-fail  up-limit
NBR list      64         1       63         0         512
Static RP     64         0       64         0          64
Anycast RP    64         0       64         0          64
timer         64         0       64         0  no-limit
prune         32         0       32         0  no-limit
pimsm J/P elem 12240       0     12240       0     48960
pimsm OIF     64         60       4         0  no-limit
mcache        64         60       4         0  no-limit
mcache hash link 997        60     937         0  no-limit
graft if no mcache 197         0     197         0  no-limit
groups        64         0       64         0  no-limit
group-memberships 64         0       64         0  no-limit
sources       256         0     256         0  no-limit
client sources 256         0     256         0  no-limit
pim/dvm intf. group 64         0       64         0  no-limit
pim/dvm global group 64         0       64         0  no-limit
MLD Resources:
  groups       64         0       64         0  no-limit
  phy-ports    64         0       64         0  no-limit
  exist-phy-port 256         0     256         0  no-limit
group-query    256         0     256         0  no-limit
group-query    256         0     256         0  no-limit
Hardware-related Resources:
HW MVID: 0 allocated for MCAST6 of total allocated 1
Total (S,G) entries 30
Total SW FWD entries 0
Total sw w/Tag MVID entries 0

```

Syntax: `show ipv6 pim [vrf vrf-name] resource`

The `vrf` parameter allows you to display IPv6 hardware resource information for the VRF instance identified by the `vrf-name` variable.

Table 41 displays the output from the `show ipv6 pim resource` command.

TABLE 41 Output from the `show ipv6 pim resource` command

Field	Description
alloc	Number of nodes of that data that are currently allocated in memory.
in-use	Number of allocated nodes in use.
avail	Number of allocated nodes are not in use.
allo-fail	Number of allocated notes that failed.
up-limit	Maximum number of nodes that can be allocated for a data structure. This may or may not be configurable, depending on the data structure

Displaying PIM traffic statistics

To display IPv6 PIM traffic statistics, enter the **show ipv6 pim traffic** command at any CLI level.

```
device# show ipv6 pim traffic
Port      Hello      Join      Prune      Assert
      [Rx      Tx]      [Rx      Tx]      [Rx      Tx]      [Rx      Tx]
MLD Statistics:
  Total Recv/Xmit 356/161
  Total Discard/chksum 0/0
```

Syntax: **show ipv6 pim [vrf vrf-name] traffic**

The **vrf** parameter allows you to display IPv6 traffic statistics for the VRF instance identified by the *vrf-name* variable.

Table 42 displays the output from the **show ipv6 pim traffic** command.

TABLE 42 Output from the **show ipv6 pim traffic** command

Field	Description
Port	The port or virtual interface on which the IPv6 PIM interface is configured.
Hello	The number of IPv6 PIM Hello messages sent or received on the interface.
J or P	The number of Join or Prune messages sent or received on the interface. NOTE Unlike PIM dense, PIM Sparse uses the same messages for Joins and Prunes.
Register	The number of Register messages sent or received on the interface.
RegStop	The number of Register Stop messages sent or received on the interface.
Assert	The number of Assert messages sent or received on the interface.
Total Recv or Xmit	The total number of IGMP messages sent and received by the device.
Total Discard or chksum	The total number of IGMP messages discarded, including a separate counter for those that failed the checksum comparison.

Clearing the IPv6 PIM forwarding cache

You can clear the IPv6 PIM forwarding cache using the **clear ipv6 pim cache** command.

```
device# clear ipv6 pim cache
```

Syntax: **clear ipv6 pim [vrf vrf-name] cache**

Use the *vrf* parameter to clear the IPv6 PIM forwarding cache for a VRF instance specified by the *vrf-name* variable.

Clearing the IPv6 PIM message counters

You can clear the IPv6 PIM message counters using the **clear ipv6 pim counters** command.

```
device# clear ipv6 pim counters
```

Syntax: **clear ipv6 pim [vrf vrf-name] counters**

Use the *vrf* parameter to clear the IPv6 PIM message counters for a VRF instance specified by the *vrf-name* variable.

Updating PIM Sparse forwarding entries with a new RP configuration

If you make changes to your static RP configuration, the entries in the IPv6 PIM Sparse multicast forwarding table continue to use the old RP configuration until they are aged out.

The **clear IPv6 pim rp-map** command allows you to update the entries in the static multicast forwarding table immediately after making RP configuration changes. This command is meant to be used with **rp-address** command.

To update the entries in an IPv6 PIM Sparse static multicast forwarding table with a new RP configuration, enter the **clear ipv6 pim rp-map** command at the privileged EXEC level of the CLI.

```
device(config)# clear ipv6 pim rp-map
```

Syntax: **clearipv6 pim [vrf vrf-name] rp-map**

Use the *vrf* parameter to clear the IPv6 PIM Sparse static multicast forwarding table for a VRF instance specified by the *vrf-name* variable.

Clearing the IPv6 PIM traffic

To clear counters on IPv6 PIM traffic, enter the **clear ipv6 pim traffic** command.

```
device# clear ipv6 pim traffic
```

Syntax: **clear ipv6 pim [vrf vrf-name] traffic**

Use the *vrf* parameter to clear counters on IPv6 PIM traffic for a VRF instance specified by the *vrf-name* variable.

Setting the maximum number of IPv6 multicast routes supported

You can use the **ipv6 max-mroute** command to define the maximum number of IPv6 multicast routes supported. The default VRF is defined using the **ipv6 max-mroute** command.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# ipv6 max-mroute
```

Syntax: **[no] ipv6 max-mroute num**

The *num* parameter specifies the maximum number of IPv6 multicast routes. If not defined by this command, the maximum value is determined by available system resources.

To define the maximum number of IPv6 multicast routes for a specified VRF, use the following commands.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# ipv6 max-mroute
```

Syntax: **[no] ipv6 router pim [vrf vrf-name]**

The *vrf* parameter specified with the **ipv6 router pim** command allows you to configure the **ipv6 max-mroute** command for a virtual routing instance (VRF) specified by the variable *vrf-name*.

Defining the maximum number of IPv6 PIM cache entries

You can use the **max-mcache** command to define the maximum number of repeated PIM traffic being sent from the same source address and being received by the same destination address. To define the maximum for the default VRF, enter the **max-mcache** command.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# max-mcache 999
```

Syntax: [no] **max-mcache** *num*

The *num* variable specifies the maximum number of IPv6 multicast cache entries for PIM in the default VRF. The maximum value and the default value that can be entered is 4K. If not defined by this command, the maximum value is determined by available system resources.

To define the maximum number of IPv6 PIM Cache entries for a specified VRF, use the following command.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# max-mcache 999
```

Syntax: [no] **ipv6 router pim** [**vrf** *vrf-name*]

The *vrf* parameter specified with the **ipv6 router pim** command allows you to configure the **max-mcache** command for a virtual routing instance (VRF) specified by the variable *vrf-name*.

Defining the maximum number of IPv6 multicast VRF CAM entries for all VRFs

You can use the following run-time command to define the maximum number of IPv6 multicast VRF CAM entries for all non-default VRF instances by entering a command such as the following.

```
device(config)# ipv6 multicast-max-all-vrf-cam 3072
```

Syntax: [no] **ipv6 multicast-max-all-vrf-cam** *cam-size*

The **ipv6 multicast-max-all-vrf-cam** command is configured at the global level. The *cam-size* variable specifies the maximum number of multicast VRF CAM entries for all non-default VRF instances. This setting does not affect the default VRF. The maximum possible value is 8000 and the default value is 2048. The maximum *cam-size* can be configured from the vrf level as well but is applied globally.

PIM Anycast RP

PIM Anycast RP is a method of providing load balancing and fast convergence to PIM RPs in an IPv6 multicast domain. The RP address of the Anycast RP is a shared address used among multiple PIM routers, known as PIM RP. The PIM RP routers create an Anycast RP set. Each router in the Anycast RP set is configured using two IP addresses: a shared RP address in their loopback address and a separate, unique IP address. The loopback address must be reachable by all PIM routers in the multicast domain. The separate, unique IP address is configured to establish static peering with other PIM routers and communication with the peers.

When the source is activated in a PIM Anycast RP domain, the PIM First Hop (FH) will register the source to the closet PIM RP. The PIM RP follows the same MSDP Anycast RP operation by decapsulating the packet and creating the (s,g) state. If there are external peers in the Anycast RP set, the router will re-encapsulate the packet with the local peering address as the source address of the encapsulation. The router will unicast the packet to all Anycast RP peers. The re-encapsulation of the data register packet to Anycast RP peers ensures source state distribution to all RPs in a multicast domain.

Configuring PIM Anycast RP

A new PIM CLI is introduced for PIM Anycast RP under the router pim submode. The PIM CLI specifies mapping of the RP and the Anycast RP peers.

To configure PIM Anycast RP, enter the following commands.

```
device(config)# ipv6 router pim
device(config-ipv6-pim-router)# rp-address 1001::1
device(config-ipv6-pim-router)# anycast-rp 1001::1 my-anycast-rp-set-acl
```

To configure PIM Anycast RP for a specified VRF, enter the commands as shown in the following example.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# rp-address 1001::1
device(config-ipv6-pim-router-vrf-blue)# anycast-rp 1001::1 my-anycast-rp-set-acl
```

Syntax: [no] anycast-rp rp-address my-anycast-rp-set-acl

The *rp address* parameter specifies a shared RP address used among multiple PIM routers.

The *my-anycast-rp-set-acl* parameter specifies a host-based simple ACL used to specify the address of the Anycast RP set, including a local address.

The following example is a configuration of PIM Anycast RP 1001:1. The example avoids using the loopback 1 interface when configuring PIM Anycast RP because the loopback 1 address could be used as a router-id. A PIM First Hop router will register the source with the closest RP. The first RP that receives the register will re-encapsulate the register to all other Anycast RP peers. Refer to [Figure 26](#) as described in the configuration of PIM Anycast RP 1001:1.

```
device(config)# interface loopback 2
device(config-lbif-2)# ipv address 1001::1/96
device(config-lbif-2)# ipv pim-sparse
device(config-lbif-2)# interface loopback 3
device(config-lbif-3)# ipv address 1:1:1::1/96
device(config-lbif-3)# ipv pim-sparse
device(config-lbif-3)# ipv6 router pim
device(config-ipv6-pim-router)# rp-address 1001::1
device(config-ipv6-pim-router)# anycast-rp 1001::1 my-anycast-rp-set
device(config-ipv6-pim-router)# ipv6 access-list my-anycast-rp-set
device(config-std-nacl)# permit ipv6 host 1:1:1::1 any
device(config-std-nacl)# permit ipv6 host 2:2:2::2 any
device(config-std-nacl)# permit ipv6 host 3:3:3::3 any
```

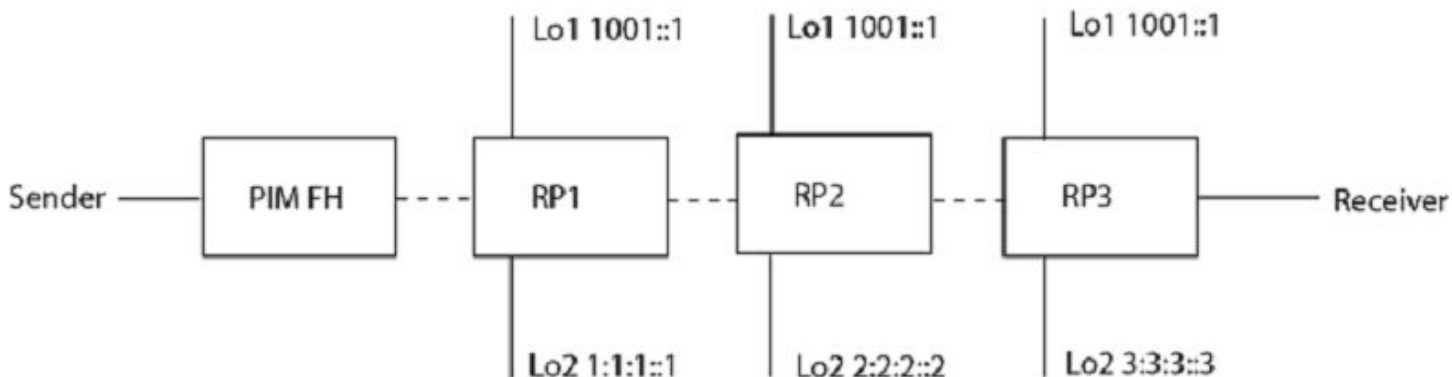
The RP shared address 1001:1 is used in the PIM domain. IP addresses 1:1:1:1, 2:2:2:2, and 3:3:3:3 are listed in the ACL that forms the self-inclusive Anycast RP set. Multiple Anycast RP instances can be configured on a system; each peer with the same or different Anycast RP set.

NOTE

The PIM Anycast CLI applies to only PIM routers running RP. All deny statements in the anycast_rp_set ACL are ignored.

The example shown in [Figure 26](#) is a PIM Anycast-enabled network with three RPs and one PIM-FH router connecting to its active source and local receiver. Loopback 2 in RP1, RP2, and RP3 each have the same IP addresses 1001:1. Loopback 3 in RP1, RP2, and RP3 each have separate IP address configured to communicate with their peers in the Anycast RP set.

FIGURE 26 Example of a PIM Anycast RP network



Displaying information for an IPv6 PIM Anycast RP interface

To display information for an IPv6 PIM Anycast RP interface, enter the **show ipv6 pim anycast-rp** command.

```
device(config)# show ipv6 pim anycast-rp
Number of Anycast RP: 1
Anycast RP: 1001::1
ACL ID: 200
ACL Name: my-anycast-rp-set
ACL Filter: SET
Peer List:
1:1:1:1
2:2:2:2
3:3:3:3
```

Syntax: `show ipv6 pim [vrf vrf-name] anycast-rp`

The *vrf* parameter allows you to display information for an IPv6 Anycast RP interface for the VRF instance identified by the *vrf-name* variable.

Table 43 describes the parameters of the **show ipv6 pim anycast-rp** command.

TABLE 43 Output from the **show ipv6 pim anycast-rp** command

Field	Description
Number of Anycast RP	Specifies the number of Anycast RP sets in the multicast domain.
Anycast RP	Specifies a shared RP address used among multiple PIM routers.
ACL ID	Specifies the ACL ID assigned.
ACL Name	Specifies the name of the Anycast RP set.
ACL Filter	Specifies the ACL filter state SET or UNSET.
Peer List	Specifies host addresses that are permitted in the Anycast RP set.

Multicast Listener Discovery and source-specific multicast protocols

Multicast Listener Discovery Version 2 (MLDv2) protocol is supported. IPv6 routers use the MLDv2 protocol to discover multicast listeners, or nodes that wish to receive multicast packets on directly attached links. MLDv2 supports source filtering, the ability of a node to send reports on traffic that is from a specific address source or from all multicast addresses except the specified address sources. The information is then provided to the source-specific multicast (SSM) routing protocols such as PIM-SSM.

The IPv6 router stores a list of multicast addresses for each attached link. For each multicast address, the IPv6 router stores a filter mode and a source list. The filter mode is set to INCLUDE if all nodes in the source list for a multicast address are in the INCLUDE state. If the filter mode is INCLUDE, then only traffic from the addresses in the source list is allowed. The filter mode is set to EXCLUDE if at least one of the nodes in the source list is in an EXCLUDE state. If the filter mode is EXCLUDE, traffic from nodes in the source list is denied and traffic from other sources is allowed.

The source list and filter mode are created when the IPv6 querier router sends a query. The querier router is the one with the lowest source IPv6 address. It sends out any of the following queries:

- **General query** - The querier sends this query to learn all multicast addresses that need to be listened to on an interface.
- **Address specific query** - The querier sends this query to determine if a specific multicast address has any listeners.
- **Address specific and source specific query** - The querier sends this query to determine if specified sources of a specific multicast address have any listeners.

In response to these queries, multicast listeners send the following reports:

- **Current state** - This report specifies the source list for a multicast address and whether the filter mode for that source list is INCLUDE or EXCLUDE.
- **Filter-mode change** - This report specifies if there has been a change to the filter mode for the source list and provides a new source list.
- **Source list change** - This report specifies the changes to the source list.

MLDv1 is compatible with IGMPv2 and MLDv2 is compatible with IGMPv3.

Enabling MLDv2

MLDv1 is enabled once PIM Sparse Mode (PIM-SM) is enabled on an interface. You then enable version 2 of MLD, the version that supports source filtering.

MLDv2 interoperates with MLDv1. MLDv1 messages are understood by MLDv2. When an IPv6 router detects that the node is operating in MLDv1 mode, the router switches to MLDv1 for that node even though queries are sent in MLDv2.

To enable IPv6 PIM-SM, enter the following command at the interface level.

```
device(config)# ipv6 router pim
device(config-if-e10000-1/1)# ipv6 pim-sparse
```

Syntax: [no] ipv6 pim-sparse

Configuring MLD parameters for default and non-default VRFs

MLD allows you to configure the following parameters on default and non-default VRFs:

- Group membership time - [Setting the group membership time](#) on page 187
- Max group address - [Defining the maximum number of MLD group addresses](#) on page 187

- Max response time - [Setting the maximum response time](#) on page 187
- Query interval - [Setting the query interval](#) on page 188
- Last listener query count - [Setting the last listener query interval](#) on page 188
- Last listener query interval - [Setting the last listener query interval](#) on page 188
- Robustness - [Setting the robustness](#) on page 189
- Version - [Setting the version](#) on page 189

Setting the group membership time

You can set the group membership time for the default VRF or for a specified VRF. Group membership time defines how long a group will remain active on an interface in the absence of a group report. Possible values are from 5 through 26,000 seconds and the default value is 260 seconds.

To define an MLD group membership time of 2000 seconds, enter the following command.

```
device(config)# ipv6 mld group-membership-time 2000
```

Syntax: [no] `ipv6 mld group-membership-time 5-26000`

To define an MLD group membership time of 2000 seconds for a specified VRF, enter the following commands.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# ipv6 mld group-membership-time 2000
```

Syntax: [no] `ipv6 router pim [vrf vrf-name]`

The *vrf* parameter specifies the virtual routing instance (VRF) specified by the variable *vrf-name*.

Defining the maximum number of MLD group addresses

You can use the following run-time command to set the maximum number of MLD addresses for the default VRF or for a specified VRF. To define this maximum for the default VRF, enter the following command.

```
device(config)# ipv6 mld max-group-address 1000
```

Syntax: [no] `ipv6 mld max-group-address num`

The *num* variable specifies the maximum number of MLD group addresses you want to make available for the default VRF. If not defined by this command, the maximum value is determined by available system resources.

To define this maximum for a specified VRF, enter the following commands.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# ipv6 mld max-group-address 1000
```

Syntax: [no] `ipv6 router pim vrf [vrf-name]`

The *vrf* parameter specifies the virtual routing instance (VRF) specified by the variable *vrf-name*.

Setting the maximum response time

You can define the maximum amount of time a multicast listener has to respond to queries by entering a command such as the following.

```
device(config)# ipv6 mld max-response-time 5
```

Syntax: [no] `ipv6 mld max-response-time seconds`

The *seconds* variable specifies the MLD maximum response time in seconds. You can specify from 1 through 10 seconds. The default is 5 seconds.

To define the maximum amount of time a multicast listener has to respond to queries for a specified VRF, enter the following commands.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# ipv6 mld max-response-time 5
```

Syntax: [no] **ipv6 router pim vrf** [*vrf-name*]

The *vrf* parameter specifies the virtual routing instance (VRF) specified by the variable *vrf-name*.

Setting the query interval

You can define the frequency at which MLD query messages are sent. For example, if you want queries to be sent every 50 seconds, enter a command such as the following.

```
device(config)# ipv6 mld query-interval 50
```

Syntax: [no] **ipv6 mld query-interval** *seconds*

The *seconds* variable specifies the MLD query interval in seconds. You can specify from 1 through 3600 seconds. The default value is 125 seconds.

To define the frequency at which MLD query messages are sent for a specified VRF, enter the following commands.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# ipv6 mld query-interval 50
```

Syntax: [no] **ipv6 router pim** [**vrf** *vrf-name*]

The *vrf* parameter specifies the virtual routing instance (VRF) specified by the variable *vrf-name*.

Setting the last listener query interval

The Last Listener Query Interval is the Maximum Response Delay inserted into Multicast-Address-Specific Queries sent in response to Done messages, and is also the amount of time between Multicast-Address-Specific Query messages. When the device receives an MLDv1 leave message or an MLDv2 state change report, it sends out a query and expects a response within the time specified by this value. Using a lower value allows members to leave groups more quickly. You can set the last listener query interval by entering a command such as the following.

```
device(config)# ipv6 mld llqi 5
```

Syntax: [no] **ipv6 mld llqi** *seconds*

The *seconds* variable sets the last listener query interval in seconds. You can specify from 1 through 10 seconds.

To set the last listener query interval for a specified VRF, enter the following commands.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# ipv6 mld llqi 5
```

Syntax: [no] **ipv6 router pim** [**vrf** *vrf-name*]

The *vrf* parameter specifies the virtual routing instance (VRF) specified by the variable *vrf-name*.

Setting the robustness

You can specify the number of times that the switch sends each MLD message from this interface. Use a higher value to ensure high reliability from MLD. You can set the robustness by entering a command such as the following.

```
device(config)# ipv6 mld robustness 3
```

Syntax: `ipv6 mld robustness seconds`

The *seconds* variable sets the MLD robustness in seconds. You can specify from 2 through 7 seconds. The default is 2 seconds.

To set the robustness for a specified VRF, enter the following commands.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# ipv6 mld robustness 3
```

Syntax: `[no] ipv6 router pim [vrf vrf-name]`

The *vrf* parameter specifies the virtual routing instance (VRF) specified by the variable *vrf-name*.

Setting the version

You can use this command to set the MLD version (1 or 2) globally. You can select the version of MLD by entering a command such as the following.

```
device(config)# ipv6 mld version 2
```

Syntax: `ipv6 mld version version-number`

The *version-number* variable sets the MLD version. You can specify 1 or 2 for the MLD version. The default version is 2.

To set the robustness for a specified VRF, enter the following commands.

```
device(config)# ipv6 router pim vrf blue
device(config-ipv6-pim-router-vrf-blue)# ipv6 mld version 2
```

Syntax: `[no] ipv6 router pim [vrf vrf-name]`

The *vrf* parameter specifies the virtual routing instance (VRF) specified by the variable *vrf-name*.

Configuring MLD parameters at the interface level

The following MLD parameters can be configured at the interface level:

- Port- version - [Specifying a port version](#) on page 189
- Static-group - [Specifying a static group](#) on page 190
- Tracking - [Enabling MLD tracking on an interface](#) on page 190
- Version - [Setting the version on an interface](#) on page 190

Specifying a port version

To set the MLD version on a virtual Ethernet interface, enter the following commands as shown in the example.

```
device(config)# interface ve 10
device(config-vif-10)# ipv6 mld port-version 2
```

Syntax: `ipv6 mld port-version version-number`

Enter 1 or 2 for *version-number*. Be sure to enter 2 if you want to use source filtering.

Specifying a static group

A multicast group is usually learned when an MLDv1 report is received. You can configure static group membership without having to receive an MLDv1 report by entering a command such as the following at the interface level.

```
device(config-if-e10000-1/1)# ipv6 mld static-group ff0d::1
```

To configure a static group membership without having to receive an MLDv1 report on a virtual Ethernet interface, enter the following commands.

```
device(config)# interface ve 10
device(config-vif-10)# ipv6 mld static-group ff0d::1
```

Syntax: `ipv6 mld static-group multicast-group-address [ethernet port-number [ethernet port-number | to port-number] *]`

Enter the IPv6 multicast group address for the *multicast-group-address*.

Enter the number of the port that will be included in this static group for the *ethernet port-number* parameter. The asterisk (*) in the syntax means that you can enter as many port numbers as you want to include in the static group. For a virtual routing interface (ve), specify the physical Ethernet ports on which to add the group address.

Enabling MLD tracking on an interface

When MLD tracking is enabled, a Layer 3 switch tracks all clients that send membership reports. When a Leave message is received from the last client, the device immediately stops forwarding to the physical port (without waiting 3 seconds to confirm that no other clients still want the traffic). To enable MLD tracking on a virtual interface, enter the following commands.

```
device(config)# interface ve 10
device(config-vif-10)# ipv6 mld tracking
```

Syntax: `ipv6 mld tracking`

Setting the version on an interface

You can use this command to set the MLD version (1 or 2) on an interface. You can select the version of MLD by entering a command such as the following.

```
device(config)# interface ve 10
device(config-vif-10)# ipv6 mld version 2
```

Syntax: `ipv6 mld version version-number`

The *version-number* variable sets the MLD version on an interface. You can specify 1 or 2 for the MLD version. The default version is 2.

Displaying MLD information

The sections below present the show commands for MLD.

Displaying MLD group information

To display the list of multicast groups, enter a command such as the following.

```
device #show ipv6 mld group
Interface e6/18 has 11 groups
  group                phy-port  static  querier  life  mode
1   ff33::6:b:1        e6/18    no      yes      0    incl
2   ff33::6:a:1        e6/18    no      yes      0    incl
3   ff33::6:9:1        e6/18    no      yes      0    incl
4   ff33::6:8:1        e6/18    no      yes      0    incl
```

```

5   ff33::6:7:1           e6/18   no     yes    0     incl
6   ff33::6:6:1           e6/18   no     yes    0     incl
7   ff33::6:5:1           e6/18   no     yes    0     incl
8   ff33::6:4:1           e6/18   no     yes    0     incl
9   ff33::6:3:1           e6/18   no     yes    0     incl
10  ff33::6:2:1            e6/18   no     yes    0     incl
11  ff33::6:1:1            e6/18   no     yes    0     incl

```

Syntax: `show ipv6 mld [vrf vrf-name] group`

The **vrf** parameter allows you to display the list of IPv6 MLD groups for the VRF instance identified by the **vrf-name** variable.

[Displaying MLD group information](#) displays the output from the `show ipv6 mld group` command.

TABLE 44 Output from the `show ipv6 mld group` command

Field	Description
Interface <i>port-number</i> has x groups	This message shows the ID of the interface and how many multicast groups it has.
#	Index for the MLD group.
group	IPv6 address of the multicast group.
phy-port	The physical port to which the group belongs.
static	Indicates if the group is a static group or not.
querier	Indicates if the multicast group is a querier or not.
life	The number of seconds the interface can remain in its current mode.
mode	Indicates if the filter mode of the multicast group is in INCLUDE or EXCLUDE.

Displaying MLD definitions for an interface

To display the MLD parameters on an interface, including the various timers, the current querying router, and whether or not MLD is enabled, enter the following command.

```

device# show ipv6 mld interface
version = 2, query int = 60, max resp time = 5, group mem time = 140
e3/1: default V2, PIM sparse, addr=fe80::20c:dbff:fe82:833a
e3/2: default V2, PIM sparse, addr=fe80::20c:dbff:fe82:833b
e6/1: default V2, PIM sparse (port down), addr=:
e6/5: default V2, PIM sparse (port down), addr=:
e6/18: default V2, PIM sparse, addr=fe80::20c:dbff:fe82:840b
      has 11 groups, Querier, default V2
      group: ff33::6:b:1, include, permit 1
      group: ff33::6:a:1, include, permit 1
      group: ff33::6:9:1, include, permit 1
      group: ff33::6:8:1, include, permit 1
      group: ff33::6:7:1, include, permit 1
      group: ff33::6:6:1, include, permit 1
      group: ff33::6:5:1, include, permit 1
      group: ff33::6:4:1, include, permit 1
      group: ff33::6:3:1, include, permit 1
      group: ff33::6:2:1, include, permit 1
      group: ff33::6:1:1, include, permit 1

```

Syntax: `show ipv6 mld [vrf vrf-name] interface [port-number]`

The **vrf** parameter allows you to display MLD parameters on an interface for the VRF instance identified by the **vrf-name** variable.

Enter a port number in the **port-number** variable if you want to display MLD information for a specific interface.

[Table 45](#) displays the output from the `show ipv6 mld interface` command.

TABLE 45 Output from the `show ipv6 mld interface` command

Field	Description
version	Version of the MLD being used.
query int	Query interval in seconds.
max resp time	Number of seconds multicast groups have to respond to queries.
group mem time	Number of seconds multicast groups can be members of this group before aging out.
(details)	The following is displayed for each interface: <ul style="list-style-type: none"> • The port ID • The default MLD version being used • The multicast protocol used • IPV6 address of the multicast interface • If the interface has groups, the group source list, IPV6 multicast address, and the filter mode are displayed.

To display the MLD parameters on an interface for a specified VRF, enter the following command as shown in the example below.

```
device(config)# show ipv mld vrf public interface
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Intf/Port|Groups| Version |Querier          | Timer  |VlRtr|Tracking
          |      | Oper   Cfg|                  |      |   |OQrr GenQ|
-----+-----+-----+-----+-----+-----+-----+-----+
v6        |    0 | 2     -            |        |    |   |Disabled
   e5/1   |    2 | 2     - fe80::20c:dbff:fee2:5000 | 11    | 0 No
v61       |    0 | 2     -            |        |    |   |Disabled
   e11/1  |    2 | 2     - Self          | 0    122 No
```

Displaying MLD settings

To display MLD settings for the "eng" VRF, enter the following command.

```
device# show ipv6 mld vrf eng settings
MLD Global Configuration
Query Interval      : 125s   Configured Interval   : 125s
Max Response Time  : 10s
Group Membership Time : 260s
Operating Version  : 2       Configured Version    : 0
Robustness Variable : 2
Last Member Query Interval: 1s   Last Member Query Count: 2
Older Host Present Timer : 260s
```

Syntax: `showipv6 mld [vrf vrf-name] settings`

The `vrf` parameter specifies that you want to display information for MLD settings for the VRF specified by the `vrf-name` variable.

[Table 46](#) displays the output from the `show ipv6 mld vrf eng settings` command.

TABLE 46 Output from the `show ipv6 mld vrf eng settings` command

Field	Description
Query Interval	How often the router will query an interface for group membership.
Configured Interval	The interval that has been configured for the router.
Max Response Time	The length of time in seconds that the router will wait for an IGMP (V1 or V2) response from an interface before concluding that the group member on that interface is down and removing it from the group.
Group Membership Time	The length of time in seconds that a group will remain active on an interface in the absence of a group report.

TABLE 46 Output from the `show ipv6 mld vrf eng settings` command (continued)

Field	Description
Operating Version	The IGMP version operating on the router.
Configured Version	The IGMP version configured on the router.
Robustness Variable	Used to fine-tune for unexpected loss on the subnet. The value is used to calculate the group interval.
Last Member Query Interval	Indicates when a leave is received; a group-specific query is sent. The last member query count is the number of queries with a time interval of (LMQT) is sent.
Last Member Query Count	Specifies the number of group-specific queries when a leave is received.

Displaying static MLD groups

The following command displays static MLD groups for the "cs" VRF.

```
device# show ipv6 mld vrf cs static
Group Address                Interface Port List
-----+-----+-----
ff1e:1::1                    v3          ethe 2/10
ff1e:a::7f                    v3          ethe 2/10
```

Syntax: `show ipv6 mld [vrf vrf-name] static`

The `vrf` parameter specifies that you want to display static MLD group information for the VRF specified by the `vrf-name` variable.

Table 47 displays the output from the `show ipv6 mld vrf cs static` command.

TABLE 47 Output from the `show ipv6 mld vrf cs static` command

Field	Description
Group Address	The address of the multicast group.
Interface Port List	The physical ports on which the multicast groups are received.

Displaying MLD traffic

To display information on MLD traffic, enter a command such as the following.

```
device# show ipv6 mld traffic
Recv  QryV1  QryV2  G-Qry  GSQry  MbrV1  MbrV2  Leave  IS_IN  IS_EX  2_IN  2_EX  ALLO  BLK
e3/1   0       0       0       0       0       0       0       0       0       0       0       0       0
e3/2   0       0       0       0       0       0       0       0       0       0       0       0       0
e6/18  0       0       0       0       0       176    0       110    0       0       0       66     0
e6/19  0       0       0       0       0       176    0       110    0       0       0       66     0
e6/20  0       0       0       0       0       176    0       110    0       0       0       66     0
e6/25  0       0       0       0       0       176    0       110    0       0       0       66     0
l1     0       0       0       0       0       0       0       0       0       0       0       0       0
Send  QryV1  QryV2  G-Qry  GSQry
e3/1   0       0       0       0
e3/2   0       0       0       0
e6/18  0       10      10      0
e6/19  0       10      10      0
e6/20  0       10      10      0
e6/25  0       10      10      0
l1     0       0       0       0
R2#
```

The report has a Receive and a Send section.

Syntax: `show ipv6 mld [vrf vrf-name] traffic`

The **vrf** parameter specifies that you want to display information on MLD traffic for the VRF specified by the *vrf-name* variable.

Table 48 displays the output from the **show ipv6 mld traffic** command.

TABLE 48 Output from the **show ipv6 mld traffic** command

Field	Description
QryV1	Number of general MLDv1 queries received or sent by the virtual routing interface.
QryV2	Number of general MLDv2 queries received or sent by the virtual routing interface.
G-Qry	Number of group-specific queries received or sent by the virtual routing interface.
GSQry	Number of source specific queries received or sent by the virtual routing interface.
MbrV1	Number of MLDv1 membership reports received.
MbrV2	Number of MLDv2 membership reports received.
Leave	Number of MLDv1 "leave" messages on the interface. (See 2_Ex for MLDv2.)
Is_IN	Number of source addresses that were included in the traffic.
Is_EX	Number of source addresses that were excluded in the traffic.
2_IN	Number of times the interface mode changed from exclude to include.
2_EX	Number of times the interface mode changed from include to exclude.
ALLOW	Number of times that additional source addresses were allowed or denied on the interface.
BLK	Number of times that sources were removed from an interface.

Clearing IPv6 MLD traffic

To clear counters on IPv6 MLD traffic, enter the following command.

```
device# clear ipv6 mld traffic
```

Syntax: **clear ipv6 mld** [**vrf** *vrf-name*] **traffic**

Use the **vrf** option to clear counters on IPv6 MLD traffic for a VRF instance specified by the *vrf-name* variable.

Clearing the IPv6 MLD group membership table cache

You can clear the IPv6 PIM group membership table cache using the following command.

```
device# clear ipv6 pim cache
```

Syntax: **clear ipv6 pim** [**vrf** *vrf-name*] **cache**

Use the **vrf** option to clear the IPv6 PIM group membership table cache for a VRF instance specified by the *vrf-name* variable.