JUNIPER NETWORKS

Multicast Protocols on the QFX Series

Release
13.2X52

Published: 2014-07-15

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Multicast Protocols on the QFX Series
13.2X52

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at http://www.juniper.net/support/eula.html. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at http://www.juniper.net/techpubs/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at http://www.juniper.net/books.

Supported Platforms

For the features described in this document, the following platforms are supported:

- QFabric System

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   system {
     scripts {
       commit {
         file ex-script.xsl;
       }
     }
   }
   interfaces {
     fxp0 {
       disable;
       unit 0 {
         family inet {
           address 10.0.0.1/24;
         }
       }
     }
   }
   ```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

   ```
   [edit]
   user@host# load merge /var/tmp/ex-script.conf
   load complete
   ```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   commit {
     file ex-script-snippet.xsl; }
   ```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

Table 1 on page xvii defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xvii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric** ***metric*****>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page at the Juniper Networks Technical Documentation site at http://www.juniper.net/techpubs/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at https://www.juniper.net/cgi-bin/docbugreport/.
- E-mail—Send your comments to **techpubs-comments@juniper.net**. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/

- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://tools.juniper.net/SerialNumberEntitlementSearch/

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see http://www.juniper.net/support/requesting-support.html.

PART 1

Overview

CHAPTER 1

Introduction to PIM Basics

PIM Overview

The predominant multicast routing protocol in use on the Internet today is Protocol Independent Multicast, or PIM. The type of PIM used on the Internet is PIM sparse mode. PIM sparse mode is so accepted that when the simple term “PIM” is used in an Internet context, some form of sparse mode operation is assumed.

PIM emerged as an algorithm to overcome the limitations of dense-mode protocols such as the Distance Vector Multicast Routing Protocol (DVMRP), which was efficient for dense clusters of multicast receivers, but did not scale well for the larger, sparser, groups encountered on the Internet. The Core Based Trees (CBT) Protocol was intended to support sparse mode as well, but CBT, with its all-powerful core approach, made placement of the core critical, and large conference-type applications (many-to-many) resulted in bottlenecks in the core. PIM was designed to avoid the dense-mode scaling issues of DVMRP and the potential performance issues of CBT at the same time.

PIM is one of the most rapidly evolving specifications on the Internet today. Since its introduction in 1995, PIM has already seen two major revisions to its packet structure (PIM version 1 [PIMv1] and PIM version 2 [PIMv2]), two major RFCs (RFC 2362 obsoleted RFC 2117), and numerous drafts describing major components of PIM, such as many-to-many trees and source-specific multicast (SSM). Long-lasting RFCs are not a feature of PIM, and virtually all of PIM must be researched, understood, and implemented directly from Internet drafts. In fact, no current RFC describes PIMv1 at all. The drafts have all expired, and PIMv1 was never issued as an official RFC.

PIM itself is not nonstandard or unstable, however. PIM has been a promising multicast routing protocol since its inception, especially PIM sparse mode, the first real sparse-mode multicast routing protocol. Work continues on PIM in a number of areas, from bidirectional trees to network management, and the rapid pace of development makes drafts essential for PIM.

PIMv1 and PIMv2 can coexist on the same routing device and even on the same interface. The main difference between PIMv1 and PIMv2 is the packet format. PIMv1 messages use Internet Group Management Protocol (IGMP) packets, whereas PIMv2 has its own IP protocol number (103) and packet structure. All routing devices connecting to an IP

subnet such as a LAN must use the same PIM version. Some PIM implementations can recognize PIMv1 packets and automatically switch the routing device interface to PIMv1. Because the difference between PIMv1 and PIMv2 involves the message format, but not the meaning of the message or how the routing device processes the PIM message, a routing device can easily mix PIMv1 and PIMv2 interfaces.

PIM is used for efficient routing to multicast groups that might span wide-area and interdomain internetworks. It is called "protocol independent" because it does not depend on a particular unicast routing protocol. Junos OS supports bidirectional mode, sparse mode, dense mode, and sparse-dense mode.

PIM operates in several modes: bidirectional mode, sparse mode, dense mode, and sparse-dense mode. In sparse-dense mode, some multicast groups are configured as dense mode (flood-and-prune, [S,G] state) and others are configured as sparse mode (explicit join to rendezvous point [RP], [*,G] state).

PIM drafts also establish a mode known as PIM source-specific mode, or PIM SSM. In PIM SSM there is only one specific source for the content of a multicast group within a given domain.

Because the PIM mode you choose determines the PIM configuration properties, you first must decide whether PIM operates in bidirectional, sparse, dense, or sparse-dense mode in your network. Each mode has distinct operating advantages in different network environments.

- In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to a downstream routing device unless the downstream routing device has sent an explicit request (by means of a join message) to the rendezvous point (RP) routing device to receive this traffic. The RP serves as the root of the shared multicast delivery tree and is responsible for forwarding multicast data from different sources to the receivers.

  Sparse mode is well suited to the Internet, where frequent interdomain join messages and prune messages are common.

- Bidirectional PIM is similar to sparse mode, and is especially suited to applications that must scale to support a large number of dispersed sources and receivers. In bidirectional PIM, routing devices build shared bidirectional trees and do not switch to a source-based tree. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Instead, it builds only group-specific (*,G) state.

- Unlike sparse mode and bidirectional mode, in which data is forwarded only to routing devices sending an explicit PIM join request, dense mode implements a *flood-and-prune* mechanism, similar to the Distance Vector Multicast Routing Protocol (DVMRP). In dense mode, a routing device receives the multicast data on the incoming interface, then forwards the traffic to the outgoing interface list. Flooding occurs periodically and is used to refresh state information, such as the source IP address and multicast group pair. If the routing device has no interested receivers for the data, and the outgoing interface list becomes empty, the routing device sends a PIM prune message upstream.

Dense mode works best in networks where few or no prunes occur. In such instances, dense mode is actually more efficient than sparse mode.

- Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.

Basic PIM Network Components

PIM dense mode requires only a multicast source and series of multicast-enabled routing devices running PIM dense mode to allow receivers to obtain multicast content. Dense mode makes sure that all multicast traffic gets everywhere by periodically flooding the network with multicast traffic, and relies on prune messages to make sure that subnets where all receivers are uninterested in that particular multicast group stop receiving packets.

PIM sparse mode is more complicated and requires the establishment of special routing devices called *rendezvous points* (*RPs*) in the network core. These routing devices are where upstream join messages from interested receivers meet downstream traffic from the source of the multicast group content. A network can have many RPs, but PIM sparse mode allows only one RP to be active for any multicast group.

If there is only one RP in a routing domain, the RP and adjacent links might become congested and form a single point of failure for all multicast traffic. Thus, multiple RPs are the rule, but the issue then becomes how other multicast routing devices find the RP that is the source of the multicast group the receiver is trying to join. This RP-to-group mapping is controlled by a special *bootstrap router* (*BSR*) running the PIM BSR mechanism. There can be more than one bootstrap router as well, also for single-point-of-failure reasons.

The bootstrap router does not have to be an RP itself, although this is a common implementation. The bootstrap router's main function is to manage the collection of RPs and allow interested receivers to find the source of their group's multicast traffic. PIM bootstrap messages are sourced from the loopback address, which is always up. The loopback address must be routable. If it is not routable, then the bootstrap router is unable to send bootstrap messages to update the RP domain members. The **show pim bootstrap** command displays only those bootstrap routers that have routable loopback addresses.

PIM SSM can be seen as a subset of a special case of PIM sparse mode and requires no specialized equipment other than that used for PIM sparse mode (and IGMP version 3).

Bidirectional PIM RPs, unlike RPs for PIM sparse mode, do not need to perform PIM Register tunneling or other specific protocol action. Bidirectional PIM RPs implement no specific functionality. RP addresses are simply a location in the network to rendezvous toward. In fact, for bidirectional PIM, RP addresses need not be loopback interface addresses or even be addresses configured on any routing device, as long as they are

covered by a subnet that is connected to a bidirectional PIM-capable routing device and advertised to the network.

Related Documentation

- *Supported IP Multicast Protocol Standards* in the *Multicast Protocols Feature Guide for Routing Devices*

PIM on Aggregated Interfaces

If you configure PIM on an aggregated (**ae-** or **as-**) interface, each of the interfaces in the aggregate is included in the multicast output interface list and carries the single stream of replicated packets in a load-sharing fashion. The multicast aggregate interface is "expanded" into its constituent interfaces in the next-hop database.

Related Documentation

- PIM Overview on page 3
- interface on page 202

CHAPTER 2

Introduction to PIM Sparse Mode

Understanding PIM Sparse Mode

A Protocol Independent Multicast (PIM) sparse-mode domain uses reverse-path forwarding (RPF) to create a path from a data source to the receiver requesting the data. When a receiver issues an explicit join request, an RPF check is triggered. A (*,G) PIM join message is sent toward the RP from the receiver's designated router (DR). (By definition, this message is actually called a join/prune message, but for clarity in this description, it is called either join or prune, depending on its context.) The join message is multicast hop by hop upstream to the ALL-PIM-ROUTERS group (224.0.0.13) by means of each router's RPF interface until it reaches the RP. The RP router receives the (*,G) PIM join message and adds the interface on which it was received to the outgoing interface list (OIL) of the rendezvous-point tree (RPT) forwarding state entry. This builds the RPT connecting the receiver with the RP. The RPT remains in effect, even if no active sources generate traffic.

NOTE: State—the (*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. S is the source IP address, G is the multicast group address, and * represents any source sending to group G. Routers keep track of the multicast forwarding state for the incoming and outgoing interfaces for each group.

When a source becomes active, the source DR encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router.

If the RP router has interested receivers in the PIM sparse-mode domain, it sends a PIM join message toward the source to build a shortest-path tree (SPT) back to the source. The source sends multicast packets out on the LAN, and the source DR encapsulates the packets in a PIM register message and forwards the message toward the RP router by means of unicast. The RP router receives PIM register messages back from the source, and thus adds a new source to the distribution tree, keeping track of sources in a PIM table. Once an RP router receives packets natively (with S,G), it sends a register stop message to stop receiving the register messages by means of unicast.

In actual application, many receivers with multiple SPTs are involved in a multicast traffic flow. To illustrate the process, we track the multicast traffic from the RP router to one receiver. In such a case, the RP router begins sending multicast packets down the RPT toward the receiver's DR for delivery to the interested receivers. When the receiver's DR receives the first packet from the RPT, the DR sends a PIM join message toward the source DR to start building an SPT back to the source. When the source DR receives the PIM join message from the receiver's DR, it starts sending traffic down all SPTs. When the first multicast packet is received by the receiver's DR, the receiver's DR sends a PIM prune message to the RP router to stop duplicate packets from being sent through the RPT. In turn, the RP router stops sending multicast packets to the receiver's DR, and sends a PIM prune message for this source over the RPT toward the source DR to halt multicast packet delivery to the RP router from that particular source.

If the RP router receives a PIM register message from an active source but has no interested receivers in the PIM sparse-mode domain, it still adds the active source into the PIM table. However, after adding the active source into the PIM table, the RP router sends a register stop message. The RP router discovers the active source's existence and no longer needs to receive advertisement of the source (which utilizes resources).

NOTE: If the number of PIM join messages exceeds the configured MTU, the messages are fragmented in IPv6 PIM sparse mode. To avoid the fragmentation of PIM join messages, the multicast traffic receives the interface MTU instead of the path MTU.

The major characteristics of PIM sparse mode are as follows:

- Routers with downstream receivers join a PIM sparse-mode tree through an explicit join message.
- PIM sparse-mode RPs are the routers where receivers meet sources.
- Senders announce their existence to one or more RPs, and receivers query RPs to find multicast sessions.
- Once receivers get content from sources through the RP, the last-hop router (the router closest to the receiver) can optionally remove the RP from the shared distribution tree (*,G) if the new source-based tree (S,G) is shorter. Receivers can then get content directly from the source.

  The transitional aspect of PIM sparse mode from shared to source-based tree is one of the major features of PIM, because it prevents overloading the RP or surrounding core links.

There are related issues regarding source, RPs, and receivers when sparse mode multicast is used:

- Sources must be able to send to all RPs.
- RPs must all know one another.
- Receivers must send explicit join messages to a known RP.

- Receivers initially need to know only one RP (they later learn about others).
- Receivers can explicitly prune themselves from a tree.
- Receivers that never transition to a source-based tree are effectively running Core Based Trees (CBT)`.

PIM sparse mode has standard features for all of these issues.

Rendezvous Point

The RP router serves as the information exchange point for the other routers. All routers in a PIM domain must provide mapping to an RP router. It is the only router that needs to know the active sources for a domain—the other routers just need to know how to reach the RP. In this way, the RP matches receivers with sources.

The RP router is downstream from the source and forms one end of the shortest-path tree. As shown in Figure 1 on page 9, the RP router is upstream from the receiver and thus forms one end of the rendezvous-point tree.

Figure 1: Rendezvous Point as Part of the RPT and SPT

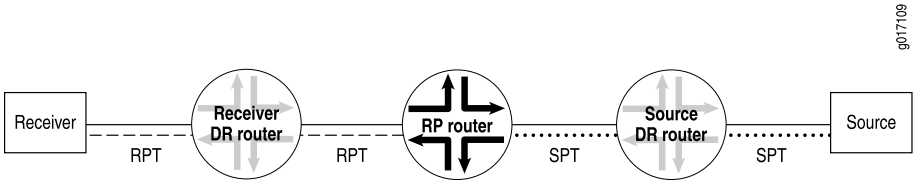

The benefit of using the RP as the information exchange point is that it reduces the amount of state in non-RP routers. No network flooding is required to provide non-RP routers information about active sources.

RP Mapping Options

RPs can be learned by one of the following mechanisms:

- Static configuration
- Anycast RP
- Auto-RP
- Bootstrap router

We recommend a static RP mapping with anycast RP and a bootstrap router (BSR) with auto-RP configuration, because static mapping provides all the benefits of a bootstrap router and auto-RP without the complexity of the full BSR and auto-RP mechanisms.

Related Documentation

- Understanding Static RP on page 11
- Understanding RP Mapping with Anycast RP on page 13
- Understanding the PIM Bootstrap Router on page 15
- *Understanding PIM Auto-RP*

Designated Router

In a PIM sparse mode (PIM-SM) domain, there are two types of designated routers to consider:

- The receiver DR sends PIM join and PIM prune messages from the receiver network toward the RP.
- The source DR sends PIM register messages from the source network to the RP.

Neighboring PIM routers multicast periodic PIM hello messages to each other every 30 seconds (the default). The PIM hello message usually includes a holdtime value for the neighbor to use, but this is not a requirement. If the PIM hello message does not include a holdtime value, a default timeout value (in Junos OS, 105 seconds) is used. On receipt of a PIM hello message, a router stores the IP address and priority for that neighbor. If the DR priorities match, the router with the highest IP address is selected as the DR.

If a DR fails, a new one is selected using the same process of comparing IP addresses.

NOTE: In PIM dense mode (PIM-DM), a DR is elected by the same process that PIM-SM uses. However, the only time that a DR has any effect in PIM-DM is when IGMPv1 is used on the interface. (IGMPv2 is the default.) In this case, the DR also functions as the IGMP Query Router because IGMPv1 does not have a Query Router election mechanism.

CHAPTER 3

Introduction to Static RP

Understanding Static RP

Protocol Independent Multicast (PIM) sparse mode is the most common multicast protocol used on the Internet. PIM sparse mode is the default mode whenever PIM is configured on any interface of the device. However, because PIM must not be configured on the network management interface, you must disable it on that interface.

Each any-source multicast (ASM) group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The rendezvous point (RP) router is the root of this shared tree and receives the multicast traffic from the source. To receive multicast traffic from the groups served by the RP, the device must determine the IP address of the RP for the source.

You can configure a static rendezvous point (RP) configuration that is similar to static routes. A static configuration has the benefit of operating in PIM version 1 or version 2. When you configure the static RP, the RP address that you select for a particular group must be consistent across all routers in a multicast domain.

One common way for the device to locate RPs is by static configuration of the IP address of the RP. A static configuration is simple and convenient. However, if the statically defined RP router becomes unreachable, there is no automatic failover to another RP router. To remedy this problem, you can use anycast RP.

Related Documentation

- Configuring Local PIM RPs on page 75
- Configuring the Static PIM RP Address on the Non-RP Routing Device on page 77

CHAPTER 4

Introduction to Anycast RP

Understanding RP Mapping with Anycast RP

Having a single active rendezvous point (RP) per multicast group is much the same as having a single server providing any service. All traffic converges on this single point, although other servers are sitting idle, and convergence is slow when the resource fails. In multicast specifically, there might be closer RPs on the shared tree, so the use of a single RP is suboptimal.

For the purposes of load balancing and redundancy, you can configure anycast RP. You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP fails, sources and receivers are taken to a new RP by means of unicast routing. When you configure anycast RP, you bypass the restriction of having one active RP per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

Anycast means that multiple RP routers share the same unicast IP address. Anycast addresses are advertised by the routing protocols. Packets sent to the anycast address are sent to the nearest RP with this address. Anycast addressing is a generic concept and is used in PIM sparse mode to add load balancing and service reliability to RPs.

Anycast RP is defined in Internet draft draft-ietf-mboned-anycast-rp-08.txt, *Anycast RP Mechanism Using PIM and MSDP*. To access Internet RFCs and drafts, go to the IETF website at http://www.ietf.org .

Related Documentation

- Configuring the Static PIM RP Address on the Non-RP Routing Device on page 77
- Example: Configuring Multiple RPs in a Domain with Anycast RP on page 85
- Example: Configuring PIM Anycast With or Without MSDP on page 79

CHAPTER 5

Introduction to PIM Bootstrap Router

Understanding the PIM Bootstrap Router

To determine which router is the rendezvous point (RP), all routers within a PIM sparse-mode domain collect bootstrap messages. A PIM sparse-mode domain is a group of routers that all share the same RP router. The domain bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routers use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

Related Documentation

- *Configuring PIM Bootstrap Properties for IPv4 or IPv6*

CHAPTER 6

Introduction to PIM Filtering

Understanding Multicast Message Filters

Multicast sources and routers generate a considerable number of control messages, especially when using PIM sparse mode. These messages form distribution trees, locate rendezvous points (RPs) and designated routers (DRs), and transition from one type of tree to another. In most cases, this multicast messaging system operates transparently and efficiently. However, in some configurations, more control over the sending and receiving of multicast control messages is necessary.

You can configure multicast filtering to control the sending and receiving of multicast control messages.

To prevent unauthorized groups and sources from registering with an RP router, you can define a routing policy to reject PIM register messages from specific groups and sources and configure the policy on the designated router or the RP router.

- If you configure the reject policy on an RP router, it rejects incoming PIM register messages from the specified groups and sources. The RP router also sends a register stop message by means of unicast to the designated router. On receiving the register stop message, the designated router sends periodic null register messages for the specified groups and sources to the RP router.
- If you configure the reject policy on a designated router, it stops sending PIM register messages for the specified groups and sources to the RP router.

NOTE: If you have configured the reject policy on an RP router, we recommend that you configure the same policy on all the RP routers in your multicast network.

NOTE: If you delete a group and source address from the reject policy configured on an RP router and commit the configuration, the RP router will register the group and source only when the designated router sends a null register message.

Related Documentation

- Filtering MAC Addresses on page 18
- Filtering RP and DR Register Messages on page 18
- Filtering MSDP SA Messages on page 40

Filtering MAC Addresses

When a router is exclusively configured with multicast protocols on an interface, multicast sets the interface media access control (MAC) filter to multicast promiscuous mode, and the number of multicast groups is unlimited. However, when the router is not exclusively used for multicasting and other protocols such as OSPF, Routing Information Protocol version 2 (RIPv2), or Network Time Protocol (NTP) are configured on an interface, each of these protocols individually requests that the interface program the MAC filter to pick up its respective multicast group only. In this case, without multicast configured on the interface, the maximum number of multicast MAC filters is limited to 20. For example, the maximum number of interface MAC filters for protocols such as OSPF (multicast group 224.0.0.5) is 20, unless a multicast protocol is also configured on the interface.

No configuration is necessary for MAC filters.

Filtering RP and DR Register Messages

You can filter Protocol Independent Multicast (PIM) register messages sent from the designated router (DR) or to the rendezvous point (RP). The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, more control over which sources an RP discovers, or which sources a DR notifies other RPs about, is desired. A high degree of control over PIM register messages is provided by RP and DR register message filtering. Message filtering also prevents unauthorized groups and sources from registering with an RP router.

Register messages that are filtered at a DR are not sent to the RP, but the sources are available to local users. Register messages that are filtered at an RP arrive from source DRs, but are ignored by the router. Sources on multicast group traffic can be limited or directed by using RP or DR register message filtering alone or together.

If the action of the register filter policy is to discard the register message, the router needs to send a register-stop message to the DR. Register-stop messages are throttled to prevent malicious users from triggering them on purpose to disrupt the routing process.

Multicast group and source information is encapsulated inside unicast IP packets. This feature allows the router to inspect the multicast group and source information before sending or accepting the PIM register message.

Incoming register messages to an RP are passed through the configured register message filtering policy before any further processing. If the register message is rejected, the RP router sends a register-stop message to the DR. When the DR receives the register-stop message, the DR stops sending register messages for the filtered groups and sources to the RP. Two fields are used for register message filtering:

- Group multicast address
- Source address

The syntax of the existing policy statements is used to configure the filtering on these two fields. The **route-filter** statement is useful for multicast group address filtering, and the **source-address-filter** statement is useful for source address filtering. In most cases, the action is to **reject** the register messages, but more complex filtering policies are possible.

Filtering cannot be performed on other header fields, such as DR address, protocol, or port. In some configurations, an RP might not send register-stop messages when the policy action is to discard the register messages. This has no effect on the operation of the feature, but the router will continue to receive register messages.

When anycast RP is configured, register messages can be sent or received by the RP. All the RPs in the anycast RP set need to be configured with the same RP register message filtering policies. Otherwise, it might be possible to circumvent the filtering policy.

Related Documentation

- Understanding RP Mapping with Anycast RP on page 13
- Configuring Register Message Filters on a PIM RP and DR on page 96

CHAPTER 7

Introduction to PIM RPT and SPT Cutover

Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees

In a shared tree, the root of the distribution tree is a router, not a host, and is located somewhere in the core of the network. In the primary sparse mode multicast routing protocol, Protocol Independent Multicast sparse mode (PIM SM), the core router at the root of the shared tree is the rendezvous point (RP). Packets from the upstream source and join messages from the downstream routers "rendezvous" at this core router.

In the RP model, other routers do not need to know the addresses of the sources for every multicast group. All they need to know is the IP address of the RP router. The RP router discovers the sources for all multicast groups.

The RP model shifts the burden of finding sources of multicast content from each router (the (S,G) notation) to the network (the (*,G) notation knows only the RP). Exactly how the RP finds the unicast IP address of the source varies, but there must be some method to determine the proper source for multicast content for a particular group.

Consider a set of multicast routers without any active multicast traffic for a certain group. When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the distribution tree for that group back to the RP, not to the actual source of the content.

To join the shared tree, or *rendezvous-point tree* (*RPT*) as it is called in PIM sparse mode, the router must do the following:

- Determine the IP address of the RP for that group. Determining the address can be as simple as static configuration in the router, or as complex as a set of nested protocols.

- Build the shared tree for that group. The router executes an RPF check on the RP address in its routing table, which produces the interface closest to the RP. The router now detects that multicast packets from this RP for this group need to flow into the router on this RPF interface.
- Send a join message out on this interface using the proper multicast protocol (probably PIM sparse mode) to inform the upstream router that it wants to join the shared tree for that group. This message is a (*,G) join message because S is not known. Only the RP is known, and the RP is not actually the source of the multicast packets. The router receiving the (*,G) join message adds the interface on which the message was received to its outgoing interface list (OIL) for the group and also performs an RPF check on the RP address. The upstream router then sends a (*,G) join message out from the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating join messages from the RPF interface, building the shared tree as it goes. The process stops when the join message reaches one of the following:

- The RP for the group that is being joined
- A router along the RPT that already has a multicast forwarding state for the group that is being joined

In either case, the branch is created, and packets can flow from the source to the RP and from the RP to the receiver. Note that there is no guarantee that the shared tree (RPT) is the shortest path tree to the source. Most likely it is not. However, there are ways to "migrate" a shared tree to an SPT once the flow of packets begins. In other words, the forwarding state can transition from (*,G) to (S,G). The formation of both types of tree depends heavily on the operation of the RPF check and the RPF table.

Related Documentation

- *Understanding Multicast Reverse Path Forwarding*

Building an RPT Between the RP and Receivers

The RPT is the path between the RP and receivers (hosts) in a multicast group (see Figure 2 on page 23). The RPT is built by means of a PIM join message from a receiver's DR:

1. A receiver sends a request to join group (G) in an Internet Group Management Protocol (IGMP) host membership report. A PIM sparse-mode router, the receiver's DR, receives the report on a directly attached subnet and creates an RPT branch for the multicast group of interest.
2. The receiver's DR sends a PIM join message to its RPF neighbor, the next-hop address in the RPF table, or the unicast routing table.
3. The PIM join message travels up the tree and is multicast to the ALL-PIM-ROUTERS group (224.0.0.13). Each router in the tree finds its RPF neighbor by using either the RPF table or the unicast routing table. This is done until the message reaches the RP

and forms the RPT. Routers along the path set up the multicast forwarding state to forward requested multicast traffic back down the RPT to the receiver.

Figure 2: Building an RPT Between the RP and the Receiver

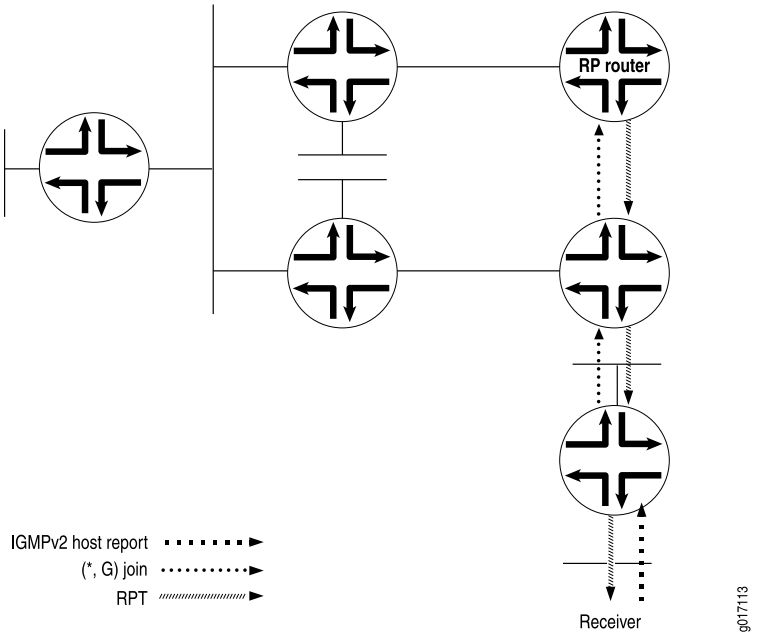

PIM Sparse Mode Source Registration

The RPT is a unidirectional tree, permitting traffic to flow down from the RP to the receiver in one direction. For multicast traffic to reach the receiver from the source, another branch of the distribution tree, called the shortest-path tree, needs to be built from the source's DR to the RP.

The shortest-path tree is created in the following way:

1. The source becomes active, sending out multicast packets on the LAN to which it is attached. The source's DR receives the packets and encapsulates them in a PIM register message, which it sends to the RP router (see Figure 3 on page 24).
2. When the RP router receives the PIM register message from the source, it sends a PIM join message back to the source.

Figure 3: PIM Register Message and PIM Join Message Exchanged

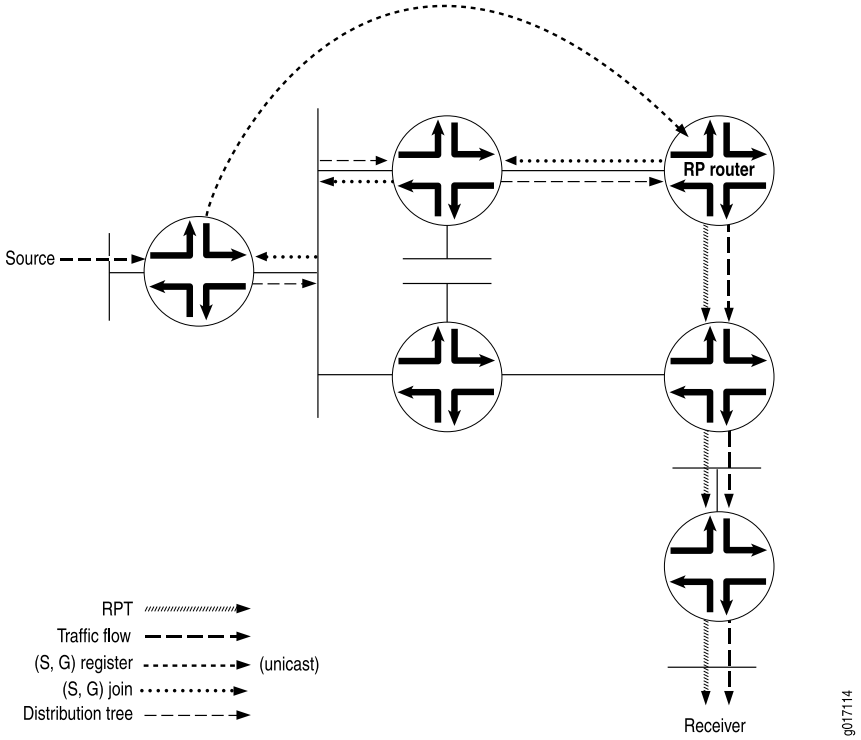

3. The source's DR receives the PIM join message and begins sending traffic down the SPT toward the RP router (see Figure 4 on page 25).

4. Once traffic is received by the RP router, it sends a register stop message to the source's DR to stop the register process.

Figure 4: Traffic Sent from the Source to the RP Router

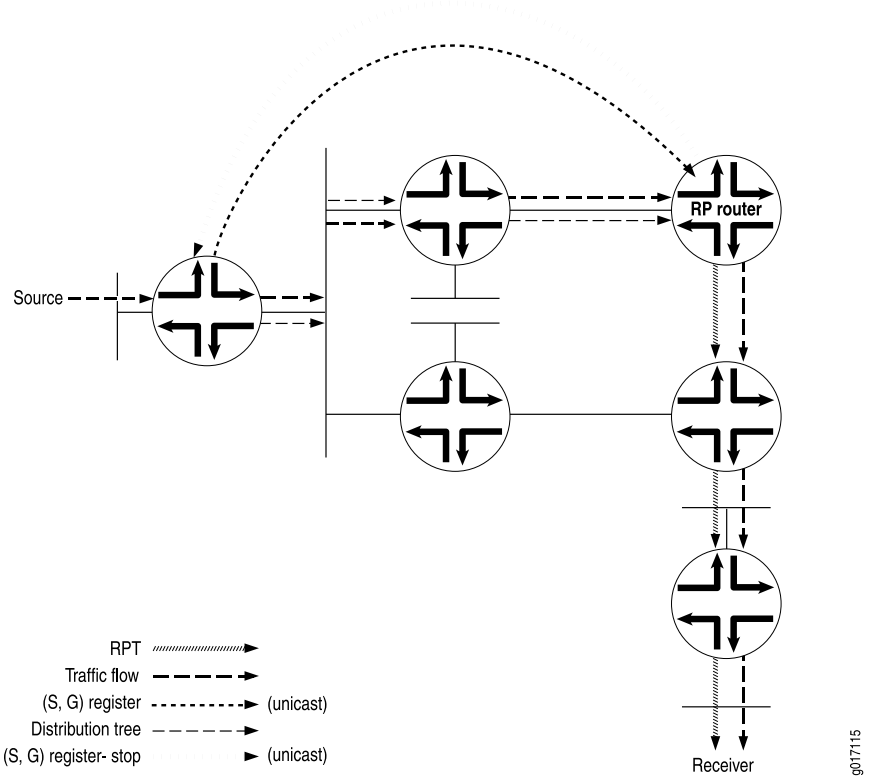

5. The RP router sends the multicast traffic down the RPT toward the receiver (see Figure 5 on page 25).

Figure 5: Traffic Sent from the RP Router Toward the Receiver

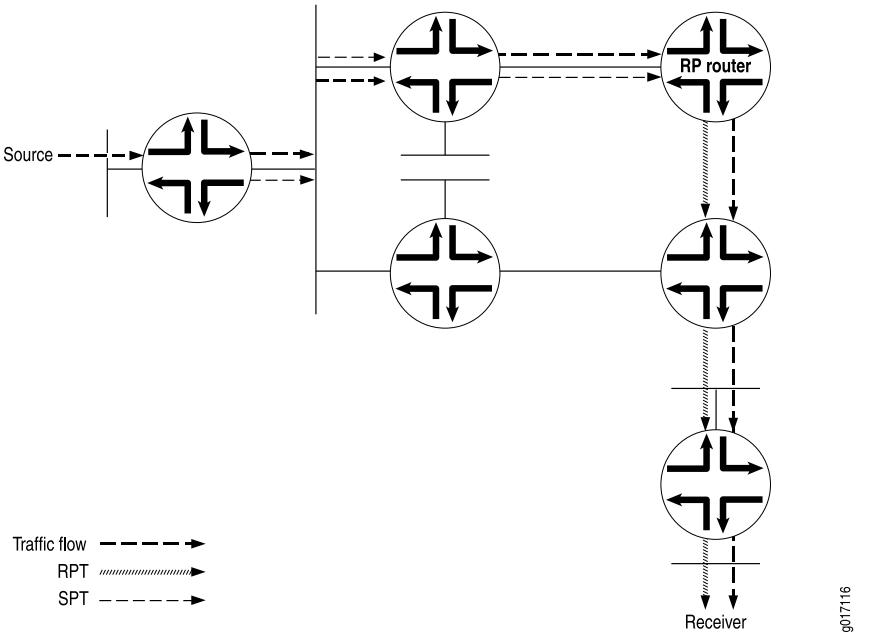

Multicast Shortest-Path Tree

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT) as well. Consider a set of multicast routers without any active multicast traffic for a certain group (that is, they have no multicast forwarding state for that group). When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the tree for that group.

To join the distribution tree, the router determines the unicast IP address of the source for that group. This address can be a simple static configuration on the router, or as complex as a set of protocols.

To build the SPT for that group, the router executes an a reverse path forwarding (RPF) check on the source address in its routing table. The RPF check produces the interface closest to the source, which is where multicast packets from this source for this group need to flow into the router.

The router next sends a join message out on this interface using the proper multicast protocol to inform the upstream router that it wants to join the distribution tree for that group. This message is an (S,G) join message because both S and G are known. The router receiving the (S,G) join message adds the interface on which the message was received to its output interface list (OIL) for the group and also performs an RPF check on the source address. The upstream router then sends an (S,G) join message out on the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating joins out on the RPF interface, building the SPT as it goes. The process stops when the join message does one of two things:

- Reaches the router directly connected to the host that is the source.
- Reaches a router that already has multicast forwarding state for this source-group pair.

In either case, the branch is created, each of the routers has multicast forwarding state for the source-group pair, and packets can flow down the distribution tree from source to receiver. The RPF check at each router makes sure that the tree is an SPT.

SPTs are always the shortest path, but they are not necessarily short. That is, sources and receivers tend to be on the periphery of a router network, not on the backbone, and multicast distribution trees have a tendency to sprawl across almost every router in the network. Because multicast traffic can overwhelm a slow interface, and one packet can easily become a hundred or a thousand on the opposite side of the backbone, it makes sense to provide a shared tree as a distribution tree so that the multicast source can be located more centrally in the network, on the backbone. This sharing of distribution trees with roots in the core network is accomplished by a multicast rendezvous point.

Related Documentation

- Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees on page 21

SPT Cutover

Instead of continuing to use the SPT to the RP and the RPT toward the receiver, a direct SPT is created between the source and the receiver in the following way:

1. Once the receiver's DR receives the first multicast packet from the source, the DR sends a PIM join message to its RPF neighbor (see Figure 6 on page 27).
2. The source's DR receives the PIM join message, and an additional (S,G) state is created to form the SPT.
3. Multicast packets from that particular source begin coming from the source's DR and flowing down the new SPT to the receiver's DR. The receiver's DR is now receiving two copies of each multicast packet sent by the source—one from the RPT and one from the new SPT.

Figure 6: Receiver DR Sends a PIM Join Message to the Source

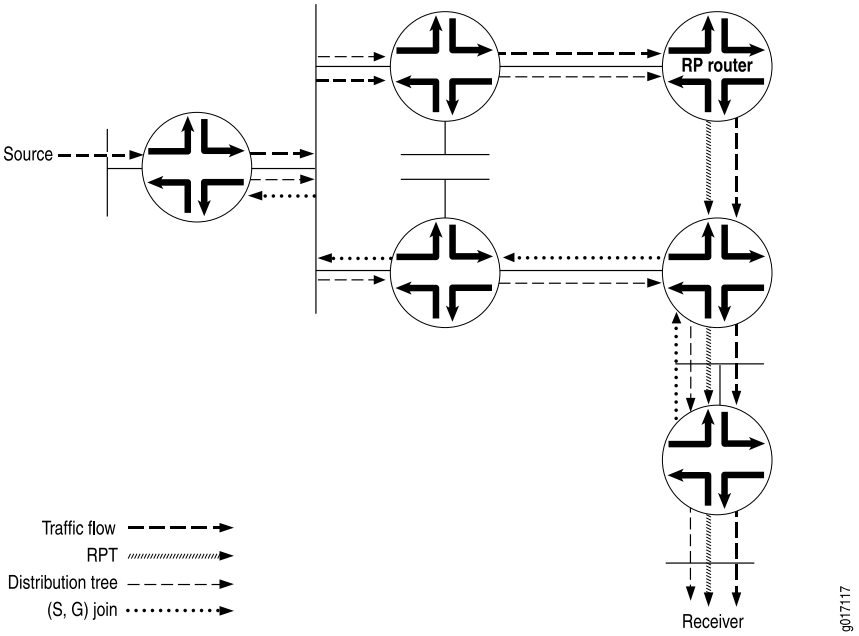

4. To stop duplicate multicast packets, the receiver's DR sends a PIM prune message toward the RP router, letting it know that the multicast packets from this particular source coming in from the RPT are no longer needed (see Figure 7 on page 28).

Figure 7: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router

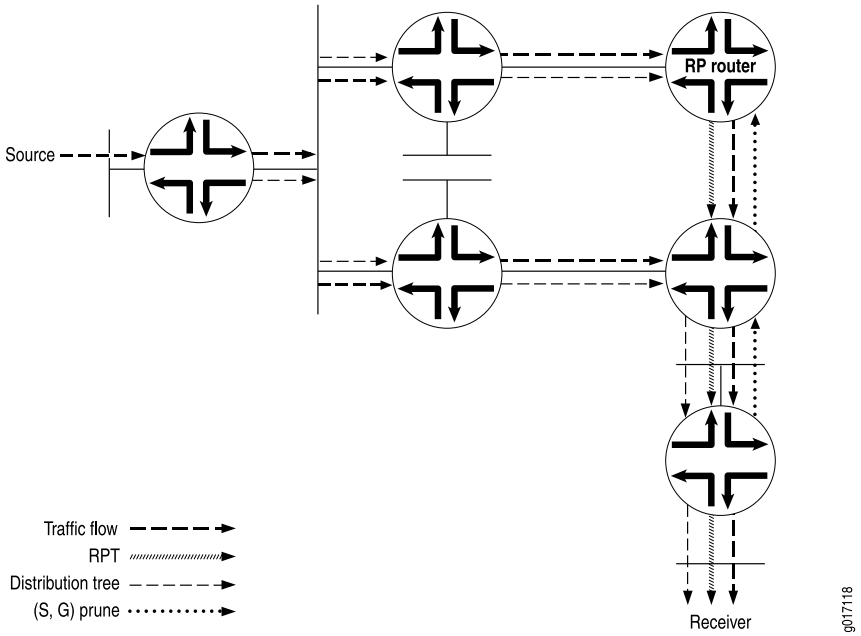

5. The PIM prune message is received by the RP router, and it stops sending multicast packets down to the receiver's DR. The receiver's DR is getting multicast packets only for this particular source over the new SPT. However, multicast packets from the source are still arriving from the source's DR toward the RP router (see Figure 8 on page 28).

Figure 8: RP Router Receives PIM Prune Message

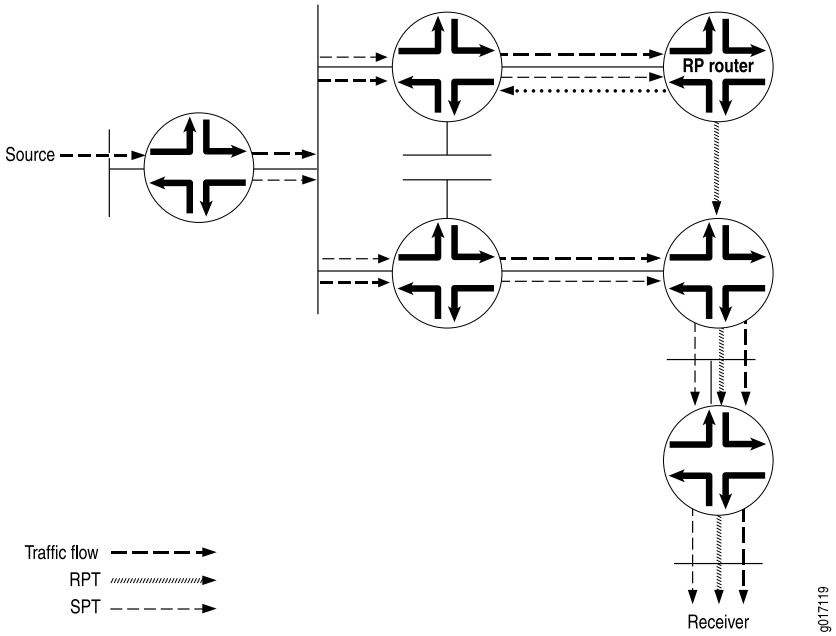

6. To stop the unneeded multicast packets from this particular source, the RP router sends a PIM prune message to the source's DR (see Figure 9 on page 29).

Figure 9: RP Router Sends a PIM Prune Message to the Source DR

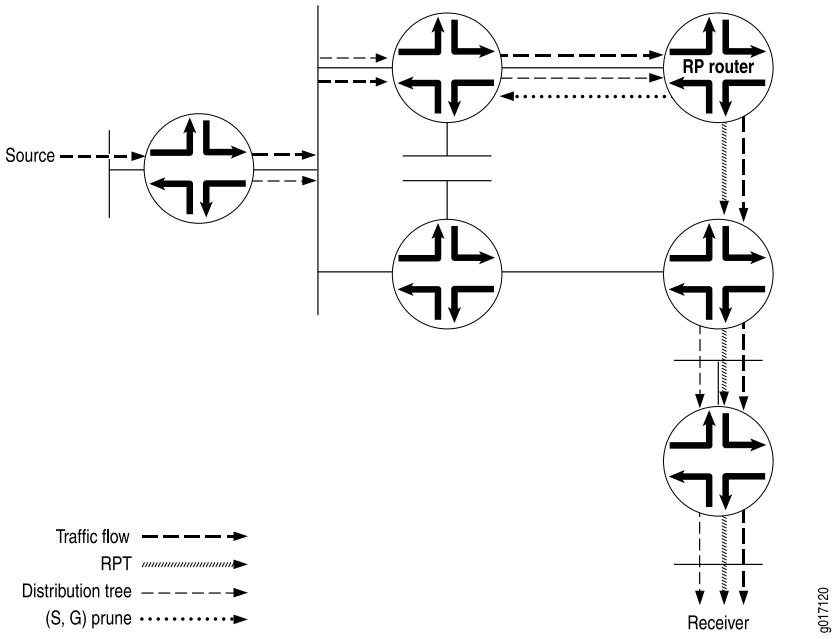

7. The receiver's DR now receives multicast packets only for the particular source from the SPT (see Figure 10 on page 29).

Figure 10: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router

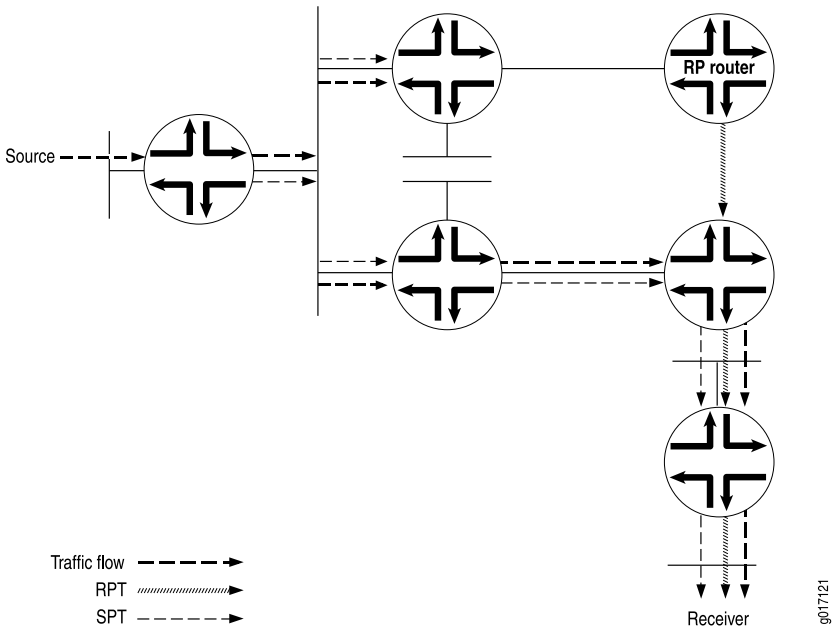

SPT Cutover Control

In some cases, the last-hop router needs to stay on the shared tree to the RP and not transition to a direct SPT to the source. You might not want the last-hop router to transition when, for example, a low-bandwidth multicast stream is forwarded from the RP to a last-hop router. All routers between last hop and source must maintain and refresh the SPT state. This can become a resource-intensive activity that does not add much to the network efficiency for a particular pair of source and multicast group addresses.

In these cases, you configure an SPT threshold policy on the last-hop router to control the transition to a direct SPT. An SPT cutover threshold of infinity applied to a source-group address pair means the last-hop router will never transition to a direct SPT. For all other source-group address pairs, the last-hop router transitions immediately to a direct SPT rooted at the source DR.

CHAPTER 8

Introduction to IGMP

Understanding Group Membership Protocols

There is a big difference between the multicast protocols used between host and router and between the multicast routers themselves. Hosts on a given subnetwork need to inform their router only whether or not they are interested in receiving packets from a certain multicast group. The source host needs to inform its routers only that it is the source of traffic for a particular multicast group. In other words, no detailed knowledge of the distribution tree is needed by any hosts; only a group membership protocol is needed to inform routers of their participation in a multicast group. Between adjacent routers, on the other hand, the multicast routing protocols must avoid loops as they build a detailed sense of the network topology and distribution tree from source to leaf. So, different multicast protocols are used for the host-router portion and the router-router portion of the multicast network.

Multicast group membership protocols enable a router to detect when a host on a directly attached subnet, typically a LAN, wants to receive traffic from a certain multicast group. Even if more than one host on the LAN wants to receive traffic for that multicast group, the router sends only one copy of each packet for that multicast group out on that interface, because of the inherent broadcast nature of LANs. When the multicast group membership protocol informs the router that there are no interested hosts on the subnet, the packets are withheld and that leaf is pruned from the distribution tree.

The Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) Protocol are the standard IP multicast group membership protocols: IGMP and MLD have several versions that are supported by hosts and routers:

- IGMPv1—The original protocol defined in RFC 1112. An explicit join message is sent to the router, but a timeout is used to determine when hosts leave a group. This process wastes processing cycles on the router, especially on older or smaller routers.
- IGMPv2—Defined in RFC 2236. Among other features, IGMPv2 adds an explicit leave message to the join message so that routers can more easily determine when a group has no interested listeners on a LAN.

- IGMPv3—Defined in RFC 3376. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or *source-specific multicast* (*SSM*).
- MLDv1—Defined in RFC 2710. MLDv1 is similar to IGMPv2.
- MLDv2—Defined in RFC 3810. MLDv2 similar to IGMPv3.

The various versions of IGMP and MLD are backward compatible. It is common for a router to run multiple versions of IGMP and MLD on LAN interfaces. Backward compatibility is achieved by dropping back to the most basic of all versions run on a LAN. For example, if one host is running IGMPv1, any router attached to the LAN running IGMPv2 can drop back to IGMPv1 operation, effectively eliminating the IGMPv2 advantages. Running multiple IGMP versions ensures that both IGMPv1 and IGMPv2 hosts find peers for their versions on the router.

Related Documentation

- *Examples: Configuring MLD*

Understanding IGMP

The IPv4 address scheme assigns class D addresses for IP multicast. IGMP is the protocol that uses these addresses, which can be in the range 224.0.0.0 to 239.255.255.255. The following addresses have specific functions or are unavailable:

- 224.0.0.0 is reserved—you cannot assign it to a group.
- 224.0.0.1 is the all-hosts address—a packet sent to this address reaches all hosts on a subnet.
- 224.0.0.2 is the all-routers address—a packet sent to this address reaches all routers on a subnet.

This implementation of IGMP complies with IGMP versions 1, 2, and 3. IGMPv3 supports source-specific join and leave messages and is backward compatible with IGMPv1 and IGMPv2.

IGMPv2 mode interfaces exchange the following types of messages between routers and hosts:

- Group membership queries
- Group membership reports
- Leave group membership messages

IGMPv3 mode interfaces exchange the following types of messages with IGMPv3 hosts:

- Group membership queries
- IGMPv3 group membership reports

IGMP manages the membership of hosts and routers in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routers. Multicast routers use IGMP to learn, for each of their attached physical networks, which groups have members.

A router receives explicit join and prune messages from those neighboring routers that have downstream group members. When PIM is the multicast protocol in use, IGMP begins the process as follows:

1. To join a multicast group, G, a host conveys its membership information through IGMP.
2. The router then forwards data packets addressed to a multicast group to only those interfaces on which explicit join messages have been received.
3. A designated router (DR) sends periodic join and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. One or more routers are automatically or statically designated as the RP, and all routers must explicitly join through the RP.
4. Each router along the path toward the RP builds a wild card (any-source) state for the group and sends join and prune messages toward the RP.

   The term *route entry* is used to refer to the state maintained in a router to represent the distribution tree.

   A route entry can include such fields as:

   - source address
   - group address
   - incoming interface from which packets are accepted
   - list of outgoing interfaces to which packets are sent
   - timers
   - flag bits

   The wild card route entry's incoming interface points toward the RP.

   The outgoing interfaces point to the neighboring downstream routers that have sent join and prune messages toward the RP as well as the directly connected hosts that have requested membership to group G.
5. This state creates a shared, RP-centered, distribution tree that reaches all group members.

IGMP is also used as the transport for several related multicast protocols (for example, Distance Vector Multicast Routing Protocol [DVMRP] and Protocol Independent Multicast version 1 [PIMv1]).

IGMP is an integral part of IP and must be enabled on all routers and hosts that need to receive IP multicast traffic.

For each attached network, a multicast router can be either a querier or a nonquerier. The querier router periodically sends general query messages to solicit group membership information. Hosts on the network that are members of a multicast group send report messages. When a host leaves a group, it sends a leave group message.

IGMP version 3 (IGMPv3) supports inclusion and exclusion lists. Inclusion lists enable you to specify which sources can send to a multicast group. This type of multicast group is called a source-specific multicast (SSM) group, and its multicast address is 232/8.

IGMPv3 provides support for source filtering. For example, a router can specify particular routers from which it accepts or rejects traffic. With IGMPv3, a multicast router can learn which sources are of interest to neighboring routers.

Exclusion mode works the opposite of an inclusion list. It allows any source but the ones listed to send to the SSM group.

IGMPv3 interoperates with versions 1 and 2 of the protocol. However, to remain compatible with older IGMP hosts and routers, IGMPv3 routers must also implement versions 1 and 2 of the protocol. IGMPv3 supports the following membership-report record types: mode is allowed, allow new sources, and block old sources.

Related Documentation

- *Supported IP Multicast Protocol Standards*

CHAPTER 9

Introduction to IGMP Snooping

IGMP Snooping Overview

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).

This IGMP snooping topic includes:

How IGMP Snooping Works

A switch usually learns unicast MAC addresses by checking the source address field of the frames it receives and then sends any traffic for that unicast address only to the appropriate interface. However, a multicast MAC address can never be the source address for a packet. As a result, when a switch receives traffic for a multicast destination address, it floods the traffic on the relevant VLAN, which can cause a significant amount of traffic to be sent unnecessarily.

IGMP snooping prevents this flooding. When you enable IGMP snooping, the switch monitors IGMP packets between receivers and multicast routers and uses the content of the packets to build a multicast cache table—a database of multicast groups and the interfaces that are connected to members of the groups. When the switch receives multicast packets, it uses the cache table to selectively forward the traffic to only the interfaces that are connected to members of the appropriate multicast groups.

NOTE: IGMP snooping is enabled by default on the default VLAN only. With versions of Junos OS for the QFX Series previous to 13.2, IGMP snooping is enabled by default on all VLANs.

NOTE: You cannot configure IGMP snooping on a secondary (private) VLAN.

How IGMP Snooping Works with Routed VLAN Interfaces

A switch can use a routed VLAN interface (RVI) to forward traffic between VLANs that connect to it. IGMP snooping works with Layer 2 interfaces and RVIs to forward multicast traffic in a switched network.

When a switch receives a multicast packet, its Packet Forwarding Engines perform a multicast lookup on the packet to determine how to forward the packet to its local interfaces. From the results of the lookup, each Packet Forwarding Engine extracts a list of Layer 3 interfaces that have ports local to the Packet Forwarding Engine. If the list includes an RVI, the switch provides a bridge multicast group ID for the RVI to the Packet Forwarding Engine.

For VLANs that include multicast receivers, the bridge multicast ID includes a sub-next-hop ID, which identifies the Layer 2 interfaces in the VLAN that are interested in receiving the multicast stream. The Packet Forwarding Engine then forwards multicast traffic to bridge multicast IDs that have multicast receivers for a given multicast group.

How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in two ways:

- By sending an unsolicited IGMP join message to a multicast router that specifies the IP multicast group that the host is attempting to join.
- By sending an IGMP join message in response to a general query from a multicast router.

A multicast router continues to forward multicast traffic to a VLAN provided that at least one host on that VLAN responds to the periodic general IGMP queries. For a host to remain a member of a multicast group, therefore, it must continue to respond to the periodic general IGMP queries.

To leave a multicast group, either a host cannot respond to the periodic general IGMP queries, which results in a “silent leave” (the only leave option for IGMPv1), or a host can send a group-specific IGMPv2 leave message.

IGMP Snooping and Forwarding Interfaces

To determine how to forward multicast traffic, a switch with IGMP snooping enabled maintains information about the following interfaces in its multicast forwarding table:

- Multicast-router interfaces—These interfaces lead toward multicast routers or IGMP queriers.
- Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups.

The switch learns about these interfaces by monitoring IGMP traffic. If an interface receives IGMP queries or Protocol Independent Multicast (PIM) updates, the switch adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the switch adds the interface to its multicast forwarding table as a group-member interface.

Table entries for interfaces that the switch learns about are subject to aging. For example, if a learned multicast-router interface does not receive IGMP queries or PIM hellos within a certain interval, the switch removes the entry for that interface from its multicast forwarding table.

NOTE: **For a switch to learn multicast-router interfaces and group-member interfaces, an IGMP querier must exist in the network. This is often a multicast router, but if there is no multicast router on the local network, you can configure the switch itself to be an IGMP querier.**

You can statically configure an interface to be a multicast-router interface or a group-member interface. The switch adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject to aging. You can have a mix of statically configured and dynamically learned interfaces on a switch.

General Forwarding Rules

Multicast traffic received on a switch interface in a VLAN on which IGMP snooping is enabled is forwarded according to the following rules.

IGMP traffic is forwarded as follows:

- IGMP general queries received on a multicast-router interface are forwarded to all other interfaces in the VLAN.
- IGMP group-specific queries received on a multicast-router interface are forwarded to only those interfaces in the VLAN that are members of the group.
- IGMP reports received on a host interface are forwarded to multicast-router interfaces in the same VLAN, but not to the other host interfaces in the VLAN.

Multicast traffic that is not IGMP traffic is forwarded as follows:

- A multicast packet with a destination address of 224.0.0.0/24 is flooded to all other interfaces on the VLAN.
- An unregistered multicast packet—that is, a packet for a group that has no current members—is forwarded to all multicast-router interfaces in the VLAN.

- A registered multicast packet is forwarded only to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN.

Using a Switch as an IGMP Querier

If IGMP snooping is enabled on a pure Layer 2 a local network (that is, Layer 3 is not enabled on the network), and there is not multicast router in the network, multicast traffic might not be properly forwarded through the network. This problem occurs if the local network is configured such that multicast traffic must be forwarded between switches in order to reach a multicast receiver. In this case, an upstream switch does not forward multicast traffic to a downstream switch (and therefore to the multicast receivers attached to the downstream switch) because the downstream switch does not forward IGMP reports to the upstream switch. You can solve this problem by configuring one of the switches to be an IGMP querier. This switch sends periodic general query packets to all the switches in the network, which ensures that the snooping membership tables are updated and prevents any multicast traffic loss.

If you configure multiple switches to be IGMP queriers, the switch with the highest (greatest) IGMP querier source address takes precedence and acts as the querier. Switches with lower IGMP querier source addresses stop sending IGMP queries unless they do not receive IGMP queries for 255 seconds. If a switch with a lower IGMP querier source address does not receive any IGMP queries during that period, it starts sending queries again.

To configure a switch to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name igmp-querier source-address source address
```

NOTE: The igmp-querier statement is not supported on QFabric systems.

Related Documentation

- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135
- Changing the IGMP Snooping Group Timeout Value on page 136
- Monitoring IGMP Snooping on page 315
- Configuring IGMP on page 113
- RFC 3171, *IANA Guidelines for IPv4 Multicast Address Assignments*
- IGMPv1—See RFC 1112, *Host extensions for IP multicasting*.
- IGMPv2—See RFC 2236, *Internet Group Management Protocol, Version 2*.
- IGMPv3—See RFC 3376, *Internet Group Management Protocol, Version 3*.

CHAPTER 10

Introduction to MSDP

Understanding MSDP

The Multicast Source Discovery Protocol (MSDP) is used to connect multicast routing domains. It typically runs on the same router as the Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP). Each MSDP router establishes adjacencies with internal and external MSDP peers similar to the way BGP establishes peers. These peer routers inform each other about active sources within the domain. When they detect active sources, the routers can send PIM sparse-mode explicit join messages to the active source.

The peer with the higher IP address passively listens to a well-known port number and waits for the side with the lower IP address to establish a Transmission Control Protocol (TCP) connection. When a PIM sparse-mode RP that is running MSDP becomes aware of a new local source, it sends source-active type, length, and values (TLVs) to its MSDP peers. When a source-active TLV is received, a peer-reverse-path-forwarding (peer-RPF) check (not the same as a multicast RPF check) is done to make sure that this peer is in the path that leads back to the originating RP. If not, the source-active TLV is dropped. This TLV is counted as a "rejected" source-active message.

The MSDP peer-RPF check is different from the normal RPF checks done by non-MSDP multicast routers. The goal of the peer-RPF check is to stop source-active messages from looping. Router R accepts source-active messages originated by Router S only from neighbor Router N or an MSDP mesh group member. For more information about configuring MSDP mesh groups, see "Example: Configuring MSDP with Active Source Limits and Mesh Groups" on page 150.

Router R locates its MSDP peer-RPF neighbor (Router N) deterministically. A series of rules is applied in a particular order to received source-active messages, and the first rule that applies determines the peer-RPF neighbor. All source-active messages from other routers are rejected.

The six rules applied to source-active messages originating at Router S received at Router R from Router X are as follows:

1. If Router X originated the source-active message (Router X is Router S), then Router X is also the peer-RPF neighbor, and its source-active messages are accepted.
2. If Router X is a member of the Router R mesh group, or is the configured peer, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
3. If Router X is the BGP next hop of the active multicast RPF route toward Router S (Router X installed the route on Router R), then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
4. If Router X is an external BGP (EBGP) or internal BGP (IBGP) peer of Router R, and the last autonomous system (AS) number in the BGP AS-path to Router S is the same as Router X's AS number, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
5. If Router X uses the same next hop as the next hop to Router S, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
6. If Router X fits none of these criteria, then Router X is not an MSDP peer-RPF neighbor, and its source-active messages are rejected.

The MSDP peers that receive source-active TLVs can be constrained by BGP reachability information. If the AS path of the network layer reachability information (NLRI) contains the receiving peer's AS number prepended second to last, the sending peer is using the receiving peer as a next hop for this source. If the split horizon information is not being received, the peer can be pruned from the source-active TLV distribution list.

Related Documentation

- Configuring MSDP on page 145

Filtering MSDP SA Messages

Along with applying MSDP source active (SA) filters on all external MSDP sessions (in and out) to prevent SAs for groups and sources from leaking in and out of the network, you need to apply bootstrap router (BSR) filters. Applying a BSR filter to the boundary of a network prevents foreign BSR messages (which announce RP addresses) from leaking into your network. Since the routers in a PIM sparse-mode domain need to know the address of only one RP router, having more than one in the network can create issues.

If you did not use multicast scoping to create boundary filters for all customer-facing interfaces, you might want to use PIM join filters. Multicast scopes prevent the actual multicast data packets from flowing in or out of an interface. PIM join filters prevent PIM sparse-mode state from being created in the first place. Since PIM join filters apply only to the PIM sparse-mode state, it might be more beneficial to use multicast scoping to filter the actual data.

NOTE: When you apply firewall filters, firewall action modifiers, such as **log**, **sample**, and **count**, work only when you apply the filter on an inbound interface. The modifiers do not work on an outbound interface.

Related Documentation

- *Understanding Multicast Administrative Scoping*
- Filtering Incoming PIM Join Messages on page 95
- Example: Configuring PIM BSR Filters on page 91

CHAPTER 11

Introduction to Source-Specific Multicast

Source-Specific Multicast Groups Overview

Source-specific multicast (SSM) is a service model that identifies session traffic by both source and group address. SSM implemented in Junos OS has the efficient explicit join procedures of Protocol Independent Multicast (PIM) sparse mode but eliminates the immediate shared tree and rendezvous point (RP) procedures using (*,G) pairs. The (*) is a wildcard referring to any source sending to group G, and “G” refers to the IP multicast group. SSM builds shortest-path trees (SPTs) directly represented by (S,G) pairs. The “S” refers to the source's unicast IP address, and the “G” refers to the specific multicast group address. The SSM (S,G) pairs are called channels to differentiate them from any-source multicast (ASM) groups. Although ASM supports both one-to-many and many-to-many communications, ASM's complexity is in its method of source discovery. For example, if you click a link in a browser, the receiver is notified about the group information, but not the source information. With SSM, the client receives both source and group information.

SSM is ideal for one-to-many multicast services such as network entertainment channels. However, many-to-many multicast services might require ASM.

To deploy SSM successfully, you need an end-to-end multicast-enabled network and applications that use an Internet Group Management Protocol version 3 (IGMPv3) or Multicast Listener Discovery version 2 (MLDv2) stack, or you need to configure SSM mapping from IGMPv1 or IGMPv2 to IGMPv3. An IGMPv3 stack provides the capability of a host operating system to use the IGMPv3 protocol. IGMPv3 is available for Windows XP, Windows Vista, and most UNIX operating systems.

SSM mapping allows operators to support an SSM network without requiring all hosts to support IGMPv3. This support exists in static (S,G) configurations, but SSM mapping also supports dynamic per-source group state information, which changes as hosts join and leave the group using IGMP.

SSM is typically supported with a subset of IGMPv3 and PIM sparse mode known as *PIM SSM*. Using SSM, a client can receive multicast traffic directly from the source. PIM SSM

uses the PIM sparse-mode functionality to create an SPT between the client and the source, but builds the SPT without the help of an RP.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through the Multicast Source Discovery Protocol (MSDP).

Understanding PIM Source-Specific Mode

RFC 1112, the original multicast RFC, supported both many-to-many and one-to-many models. These came to be known collectively as any-source multicast (ASM) because ASM allowed one or many sources for a multicast group's traffic. However, an ASM network must be able to determine the locations of all sources for a particular multicast group whenever there are interested listeners, no matter where the sources might be located in the network. In ASM, the key function of *source discovery* is a required function of the network itself.

Multicast source discovery appears to be an easy process, but in sparse mode it is not. In dense mode, it is simple enough to flood traffic to every router in the whole network so that every router learns the source address of the content for that multicast group. However, the flooding presents scalability and network resource use issues and is not a viable option in sparse mode.

PIM sparse mode (like any sparse mode protocol) achieves the required source discovery functionality without flooding at the cost of a considerable amount of complexity. The RP routers must be added and must know all multicast sources, and complicated shared distribution trees must be built to the RPs.

In an environment where many sources come and go, such as for a videoconferencing service, ASM is appropriate. However, by ignoring the many-to-many model and focusing attention on the one-to-many source-specific multicast (SSM) model, several commercially promising multicast applications, such as television channel distribution over the Internet, might be brought to the Internet much more quickly and efficiently than if full ASM functionality were required of the network.

PIM SSM is simpler than PIM sparse mode because only the one-to-many model is supported. Initial commercial multicast Internet applications are likely to be available to *subscribers* (that is, receivers that issue join messages) from only a single source (a special case of SSM covers the need for a backup source). PIM SSM therefore forms a subset of PIM sparse mode. PIM SSM builds shortest-path trees (SPTs) rooted at the source immediately because in SSM, the router closest to the interested receiver host is informed of the unicast IP address of the source for the multicast traffic. That is, PIM SSM bypasses the RP connection stage through shared distribution trees, as in PIM sparse mode, and goes directly to the source-based distribution tree.

PIM SSM introduces new terms for many of the concepts in PIM sparse mode. PIM SSM can technically be used in the entire 224/4 multicast address range, although PIM SSM operation is guaranteed only in the 232/8 range (232.0.0/24 is reserved). The new SSM terms are appropriate for Internet video applications and are summarized in Table 3 on page 45.

Table 3: ASM and SSM Terminology

Term	Any-Source Multicast	Source-Specific Multicast
Address identifier	G	S,G
Address designation	group	channel
Receiver operations	join, leave	subscribe, unsubscribe
Group address range	224/4 excluding 232/8	224/4 (guaranteed only for 232/8)

Although PIM SSM describes receiver operations as *subscribe* and *unsubscribe*, the same PIM sparse mode join and leave messages are used by both forms of the protocol. The terminology change distinguishes ASM from SSM even though the receiver messages are identical.

PIM SSM

PIM source-specific multicast (SSM) uses a subset of PIM sparse mode and IGMP version 3 (IGMPv3) to allow a client to receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the receiver and the source, but builds the SPT without the help of an RP.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the **ssm-groups** statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.

You can also configure the Junos OS to accept any-source multicast (ASM) join messages (*,G) for group addresses that are within the default or configured range of source-specific multicast (SSM) groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through MSDP.

Deploying SSM is easy. You need to configure PIM sparse mode on all router interfaces and issue the necessary SSM commands, including specifying IGMPv3 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group member interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3, are used in PIM SSM. As sources become active and start sending multicast packets, interested receivers in the SSM group receive the multicast packets.

In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3), announcing a desire to join group G and source S (see Figure 11 on page 46).

The directly connected PIM sparse-mode router, the receiver's DR, sends an (S,G) join message to its RPF neighbor for the source. Notice in Figure 11 on page 46 that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 11: Receiver Announces Desire to Join Group G and Source S

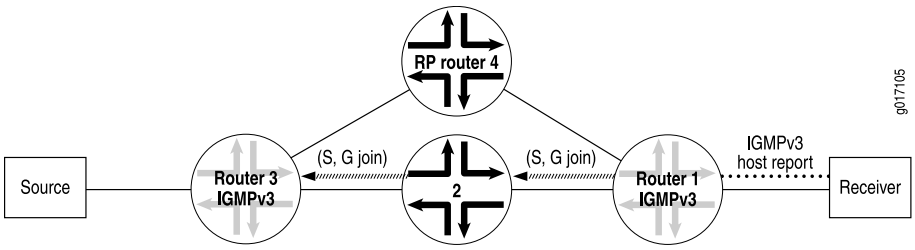

The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In Figure 12 on page 46, the source tree is built across the network to Router 3, the last-hop router connected to the source.

Figure 12: Router 3 (Last-Hop Router) Joins the Source Tree

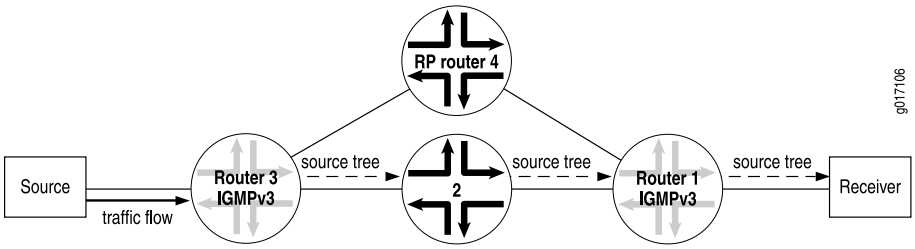

Using the source tree, multicast traffic is delivered to the subscribing host (see Figure 13 on page 46).

Figure 13: (S,G) State Is Built Between the Source and the Receiver

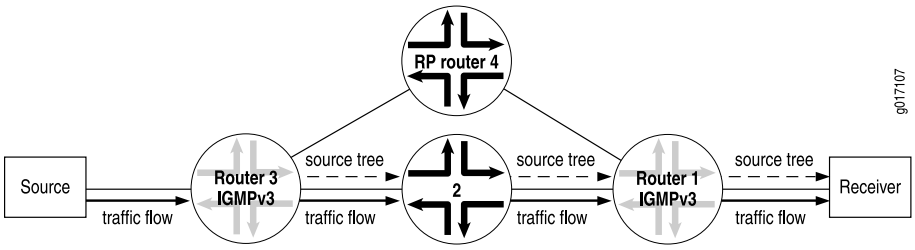

To configure additional SSM groups, include the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level.

Related Documentation

- Source-Specific Multicast Groups Overview on page 43
- Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 165

CHAPTER 12

Introduction to Multicast VLAN Registration

Understanding Multicast VLAN Registration

Multicast VLAN registration (MVR) enables you to efficiently distribute IPTV multicast streams across an Ethernet ring-based Layer 2 network and reduce the amount of bandwidth consumed by this multicast traffic.

In a standard Layer 2 network, a multicast stream received on one VLAN is never distributed to interfaces outside that VLAN. If hosts in multiple VLANs request the same multicast stream, a separate copy of that multicast stream is distributed to the requesting VLANs.

MVR introduces the concept of a *multicast source VLAN* (MVLAN), which is created by MVR and becomes the only VLAN over which IPTV multicast traffic flows throughout the Layer 2 network. The Juniper Networks EX Series Ethernet Switch or the QFX Series that is enabled for MVR selectively forwards IPTV multicast traffic from interfaces on the MVLAN (source interfaces) to hosts that are connected to interfaces that are not part of the MVLAN. These interfaces are known as *MVR receiver ports*. The MVR receiver ports can receive traffic from a port on the MVLAN but cannot send traffic onto the MVLAN, and they remain in their own VLANs for bandwidth and security reasons.

This topic includes:

How MVR Works

In many ways, MVR is similar to IGMP snooping. Both MVR and IGMP snooping monitor IGMP join and leave messages and build forwarding tables based on the media access control (MAC) addresses of the hosts sending those IGMP messages. Whereas IGMP snooping operates within a given VLAN to regulate multicast traffic, MVR can operate with hosts on different VLANs in a Layer 2 network to selectively deliver IPTV multicast traffic to requesting hosts, thereby reducing the amount of bandwidth needed to forward multicast traffic.

When you configure an MVLAN, you assign a range of multicast group addresses to it. You then configure other VLANs to be MVR receiver VLANs, which receive multicast streams from the MVLAN. The MVR receiver ports comprise all the interfaces that exist on any of the MVR receiver VLANs. Interfaces that are on the MVLAN itself cannot be MVR receiver ports for that MVLAN.

NOTE: MVR is supported on VLANs running IGMP version 2 (IGMPv2) only.

MVR Modes

MVR operates in two modes: MVR transparent mode and MVR proxy mode. Both modes enable MVR to forward only one copy of a multicast stream to the Layer 2 network.

- MVR Transparent Mode on page 48
- MVR Proxy Mode on page 48

MVR Transparent Mode

In MVR transparent mode, the switch receives one copy of each IPTV multicast stream and then replicates the stream only to those hosts that want to receive it, while forwarding all other types of multicast traffic without modification. Transparent mode is the default mode.

The switch handles IGMP packets destined for both the multicast source VLAN and multicast receiver VLANs in the same way that it handles them when MVR is not being used. That is, when a host on a VLAN sends IGMP join and leave messages, the switch floods the messages to all router interfaces in the VLAN. Similarly, when a VLAN receives IGMP queries from its router interfaces, it floods the queries to all interfaces in the VLAN.

If a host on a multicast receiver port joins an MVR group on the multicast receiver VLAN, the appropriate bridging entry is added and the MVLAN forwards that group's IPTV multicast traffic on that port (even though that port is not in the MVLAN). Likewise, if a host on a multicast receiver port leaves an MVR group on the multicast receiver VLAN, the appropriate bridging entry is deleted, and the MVLAN stops forwarding that group's IPTV multicast traffic on that port. In addition, you can configure the switch to statically install the bridging entries on the multicast receiver VLAN.

MVR Proxy Mode

When you use MVR in proxy mode, the switch acts as a proxy for any MVR group in both the upstream and downstream directions. In the downstream direction, the switch acts as the querier for the groups in the MVR receiver VLANs. In the upstream direction, the switch originates the IGMP reports and leaves and answers IGMP queries from multicast routers. When the MVR receiver VLANs receive IGMP joins and leaves, the switch creates bridging entries on the MVLAN as needed, as it does in MVR transparent mode. In addition, the switch sends out IGMP joins and leaves on the MVLAN based on these bridging entries.

Configuring MVR proxy mode on the MVLAN automatically enables IGMP snooping proxy mode on all MVR receiver VLANs as well as on the MVLAN.

Related Documentation

- *Understanding FIP Snooping, FBF, and MVR Filter Scalability*
- Example: Configuring Multicast VLAN Registration on page 140
- Configuring Multicast VLAN Registration (CLI Procedure) on page 139

PART 2

Configuration

CHAPTER 13

Optimizing Multicast Flows on QFabric Systems

Optimizing the Number of Multicast Flows on QFabric Systems

Because of the distributed nature of QFabric systems, the default configuration does not allow the maximum number of supported Layer 3 multicast flows to be created. To allow a QFabric system to create the maximum number of supported flows, configure the following statement:

```
set fabric routing-options multicast fabric-optimized-distribution
```

After configuring this statement, you must reboot the QFabric Director group to make the change take effect.

Related Documentation

-

CHAPTER 14

PIM Basics

Changing the PIM Version

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default for rendezvous point (RP) mode (at the **[edit protocols pim rp static address *address*]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface *interface-name*]** hierarchy level). Explicitly configured versions override the defaults.

To configure the PIM version, include the **version** statement:

```
version (1 | 2);
```

Modifying the PIM Hello Interval

Routing devices send hello messages at a fixed interval on all PIM-enabled interfaces. By using hello messages, routing devices advertise their existence as PIM routing devices on the subnet. With all PIM-enabled routing devices advertised, a single designated router for the subnet is established.

When a routing device is configured for PIM, it sends a hello message at a 30-second default interval. The interval range is from 0 through 255. When the interval counts down to 0, the routing device sends another hello message, and the timer is reset. A routing device that receives no response from a neighbor in 3.5 times the interval value drops the neighbor. In the case of a 30-second interval, the amount of time a routing device waits for a response is 105 seconds.

If a PIM hello message contains the hold-time option, the neighbor timeout is set to the hold-time sent in the message. If a PIM hello message does not contain the hold-time option, the neighbor timeout is set to the default hello hold time.

To modify how often the routing device sends hello messages out of an interface:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

   ```
   [edit routing-instances PIM.master protocols pim interface fe-3/0/2.0]
   user@host# set hello-interval 255
   ```

2. Verify the configuration by checking the **Hello Option Holdtime** field in the output of the **show pim neighbors detail** command.

   ```
   user@host> show pim neighbors detail
   Instance: PIM.master
   Interface: fe-3/0/2.0
   Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
   Hello Option Holdtime: 255 seconds
   Hello Option DR Priority: 1
   Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
   Join Suppression supported
   Rx Join: Group Source Timeout
   225.1.1.1 192.168.195.78 0
   225.1.1.1 0

   Interface: lo0.0
   Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
   Hello Option Holdtime: 255 seconds
   Hello Option DR Priority: 1
   Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
   Join Suppression supported

   Interface: pd-6/0/0.32768
   Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
   Hello Option Holdtime: 255 seconds
   Hello Option DR Priority: 0
   Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
   Join Suppression supported
   ```

Related Documentation

- show pim neighbors on page 434 in the CLI Explorer

Preserving Multicast Performance by Disabling Response to the ping Utility

The ping utility uses ICMP Echo messages to verify connectivity to any device with an IP address. However, in the case of multicast applications, a single ping sent to a multicast address can degrade the performance of routers because the stream of packets is replicated multiple times.

You can disable the router's response to ping (ICMP Echo) packets sent to multicast addresses. The system responds normally to unicast ping packets.

To disable the router's response to ping packets sent to multicast addresses:

1. Include the **no-multicast-echo** statement:

   ```
   [edit system]
   user@host# set no-multicast-echo
   ```

2. Verify the configuration by checking the **echo drops with broadcast or multicast destination address** field in the output of the **show system statistics icmp** command.

   user@host> **show system statistics icmp**

```
icmp:
0 drops due to rate limit
0 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
echo reply: 21
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
100 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
echo: 21
21 message responses generated
```

Related Documentation

- *Configuring the Junos OS to Disable the Routing Engine Response to Multicast Ping Packets* in the *Junos OS Administration Library for Routing Devices*
- *show system statistics icmp* in the CLI Explorer

Configuring PIM Trace Options

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
assert	Trace assert messages, which are used to resolve which of the parallel routers connected to a multiaccess LAN is responsible for forwarding packets to the LAN.
autorp	Trace bootstrap, RP, and auto-RP messages.
bidirectional-df-election	Trace bidirectional PIM designated-forwarder (DF) election events.
bootstrap	Trace bootstrap messages, which are sent periodically by the PIM domain's bootstrap router and are forwarded, hop by hop, to all routers in that domain.
general	Trace general events.
graft	Trace graft and graft acknowledgment messages.

Flag	Description
hello	Trace hello packets, which are sent so that neighboring routers can discover one another.
join	Trace join messages, which are sent to join a branch onto the multicast distribution tree.
mdt	Trace messages related to multicast data tunnels.
normal	Trace normal events.
nsr-synchronization	Trace nonstop routing synchronization events
packets	Trace all PIM packets.
policy	Trace poison-route-reverse packets.
prune	Trace prune messages, which are sent to prune a branch off the multicast distribution tree.
register	Trace register and register-stop messages. Register messages are sent to the RP when a multicast source first starts sending to a group.
route	Trace routing information.
rp	Trace candidate RP advertisements.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on PIM packets of a particular type.

To configure tracing operations for PIM:

1. (Optional) Configure tracing at the [**routing-options** hierarchy level to trace all protocol packets.

   ```
   [edit routing-options traceoptions]
   user@host# set file all-packets-trace
   user@host# set flag all
   ```

2. Configure the filename for the PIM trace file.

   ```
   [edit protocols pim traceoptions]
   user@host# set file pim-trace
   ```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols pim traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols pim traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols pim traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags.
   Suppose you are troubleshooting issues with PIM version 1 control packets that are received on an interface configured for PIM version 2. The following example shows how to trace messages associated with this problem.

```
[edit protocols pim traceoptions]
user@host# set flag packets | match “Rx V1 Require V2”
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/pim-trace
```

Related Documentation

- PIM Overview on page 3
- *Junos OS Tracing and Logging Operations* in the *Junos OS Administration Library for Routing Devices*

Disabling PIM

By default, when configured, the PIM protocol is enabled on all interfaces for all families. If desired, you can disable PIM at the protocol, interface, or family hierarchy levels.

The hierarchy in which you configure PIM is critical. In general, the most specific configuration takes precedence. However, if PIM is disabled at the protocol level, then any disable statements with respect to an interface or family are ignored.

For example, the order of precedence for disabling PIM on a particular interface family is:

1. If PIM is disabled at the **[edit protocols pim interface *interface-name* family]** hierarchy level, then PIM is disabled for that interface family.
2. If PIM is not configured at the **[edit protocols pim interface *interface-name* family]** hierarchy level, but is disabled at the **[edit protocols pim interface *interface-name*]** hierarchy level, then PIM is disabled for all families on the specified interface.
3. If PIM is not configured at either the **[edit protocols pim interface *interface-name* family]** hierarchy level or the **[edit protocols pim interface *interface-name*]** hierarchy level, but is disabled at the **[edit protocols pim]** hierarchy level, then the PIM protocol is disabled globally for all interfaces and all families.

The following sections describe how to disable PIM at the various hierarchy levels.

Disabling the PIM Protocol

You can explicitly disable the PIM protocol. Disabling the PIM protocol disables the protocol for all interfaces and all families. This is accomplished at the **[edit protocols pim]** hierarchy level:

```
[edit protocols]
pim {
   disable;
}
```

To disable the PIM protocol:

1. Include the **disable** statement.

   ```
   user@host# set protocols pim disable
   ```
2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

   ```
   user@host# run show protocols pim
   ```

Disabling PIM On an Interface

You can disable the PIM protocol on a per-interface basis. This is accomplished at the **[edit protocols pim interface *interface-name*]** hierarchy level:

```
[edit protocols]
pim {
   interface interface-name {
      disable;
   }
}
```

To disable PIM on an interface:

1. Include the **disable** statement.

   ```
   user@host# set protocols pim interface fe-0/1/0 disable
   ```
2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

   ```
   user@host# run show protocols pim
   ```

Disabling PIM for a Family

You can disable the PIM protocol on a per-family basis. This is accomplished at the **[edit protocols pim family]** hierarchy level:

```
[edit protocols]
pim {
   family inet {
      disable;
   }
   family inet6 {
      disable;
   }
}
```

To disable PIM for a family:

1. Include the **disable** statement.

   ```
   user@host# set protocols pim family inet disable
   user@host# set protocols pim family inet6 disable
   ```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

   ```
   user@host# run show protocols pim
   ```

Disabling PIM for a Rendezvous Point

You can disable the PIM protocol for a rendezvous point (RP) on a per-family basis. This is accomplished at the **[edit protocols pim rp local family]** hierarchy level:

```
[edit protocols]
pim {
   rp {
      local {
         family inet {
            disable;
         }
         family inet6 {
            disable;
         }
      }
   }
}
```

To disable PIM for an RP family:

1. Use the **disable** statement.

   ```
   user@host# set protocols pim rp local family inet disable
   user@host# set protocols pim rp local family inet6 disable
   ```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

   ```
   user@host# run show protocols pim
   ```

CHAPTER 15

PIM Designated Router

Configuring Interface Priority for PIM Designated Router Selection

A designated router (DR) sends periodic join messages and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. When a Protocol Independent Multicast (PIM) router learns about a source, it originates a Multicast Source Discovery Protocol (MSDP) source-address message if it is the DR on the upstream interface.

By default, every PIM interface has an equal probability (priority 1) of being selected as the DR. Configuring the interface DR priority helps ensure that changing an IP address does not alter your forwarding model.

To configure the interface designated router priority:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

   ```
   [edit routing-instances PIM.master protocols pim interface ge-0/0/0.0 family inet]
   user@host# set priority 5
   ```

2. Verify the configuration by checking the **Hello Option DR Priority** field in the output of the **show pim neighbors detail** command.

   ```
   user@host> show pim neighbors detail
   ```

   ```
   Instance: PIM.master
   Interface: ge-0/0/0.0
   Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
   Hello Option Holdtime: 65535 seconds
   Hello Option DR Priority: 5
   Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
   Join Suppression supported
   Rx Join: Group Source Timeout
   225.1.1.1 192.168.195.78 0
   225.1.1.1 0

   Interface: lo0.0
   Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
   Hello Option Holdtime: 65535 seconds
   Hello Option DR Priority: 1
   ```

```
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

Related Documentation

- Configuring PIM Designated Router Election on Point-to-Point Links on page 64
- Understanding PIM Sparse Mode on page 7
- **show pim neighbors on page 434** in the CLI Explorer

Configuring PIM Designated Router Election on Point-to-Point Links

To comply with the latest PIM drafts, enable designated router (DR) election on all PIM interfaces, including point-to-point (P2P) interfaces. (DR election is enabled by default on all other interfaces.) One of the two routers might join a multicast group on its P2P link interface. The DR on that link is responsible for initiating the relevant join messages.

To enable DR election on point-to-point interfaces:

1. On both point-to-point link routers, configure the router globally or in the routing instance. This example shows the configuration for the routing instance.

   ```
   [edit routing-instances PIM.master protocols pim]
   user@host# set dr-election-on-p2p
   ```

2. Verify the configuration by checking the **State** field in the output of the **show pim interfaces** command. The possible values for the **State** field are DR, NotDR, and P2P. When a point-to-point link interface is elected to be the DR, the interface state becomes DR instead of P2P.

3. If the **show pim interfaces** command continues to report the P2P state, consider running the **restart routing** command on both routers on the point-to-point link. Then recheck the state.

   

   CAUTION: Do not restart a software process unless specifically asked to do so by your Juniper Networks customer support representative. Restarting a software process during normal operation of a routing platform could cause interruption of packet forwarding and loss of data.

   ```
   [edit]
   user@host# run restart routing
   ```

Related Documentation

- Understanding PIM Sparse Mode on page 7
- Configuring Interface Priority for PIM Designated Router Selection on page 63
- **show pim interfaces on page 417** in the CLI Explorer

CHAPTER 16

PIM Sparse Mode

Enabling PIM Sparse Mode

In PIM sparse mode (PIM-SM), the assumption is that very few of the possible receivers want packets from a source, so the network establishes and sends packets only on branches that have at least one leaf indicating (by message) a desire for the traffic. WANs are appropriate networks for sparse-mode operation.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default. You do not need to configure Internet Group Management Protocol (IGMP) version 2 for a sparse mode configuration. After you enable PIM, by default, IGMP version 2 is also enabled.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default for rendezvous point (RP) mode (at the **[edit protocols pim rp static address *address*]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface *interface-name*]** hierarchy level). Explicitly configured versions override the defaults. The following example explicitly configures PIMv2 on the interfaces.

You can configure PIM sparse mode globally or for a routing instance. This example shows how to configure PIM sparse mode globally on all interfaces. It also shows how to configure a static RP router and how to configure the non-RP routers.

To configure the router properties for PIM sparse mode:

1. Configure the static RP router.

   ```
   [edit protocols pim]
   user@host# set rp local family inet address 192.168.3.253
   ```

2. Configure the RP router interfaces. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
[edit protocols pim]
user@host# set interface all mode sparse
user@host# set interface all version 2
user@host# set interface fxp0.0 disable
```

3. Configure the non-RP routers. Include the following configuration on all of the non-RP routers.

```
[edit protocols pim]
user@host# set rp static address 192.168.3.253 version 2
user@host# set interface all mode sparse
user@host# set interface all version 2
user@host# set interface fxp0.0 disable
```

4. Monitor the operation of PIM sparse mode.
   - show pim interfaces
   - show pim join
   - show pim neighbors
   - show pim rps

Related Documentation
- Understanding PIM Sparse Mode on page 7

Configuring PIM Join Load Balancing

By default, PIM join messages are sent toward a source based on the RPF routing table check. If there is more than one equal-cost path toward the source, then one upstream interface is chosen to send the join message. This interface is also used for all downstream traffic, so even though there are alternative interfaces available, the multicast load is concentrated on one upstream interface and routing device.

For PIM sparse mode, you can configure PIM join load balancing to spread join messages and traffic across equal-cost upstream paths (interfaces and routing devices) provided by unicast routing toward a source. PIM join load balancing is only supported for PIM sparse mode configurations.

PIM join load balancing is supported on draft-rosen multicast VPNs (also referred to as dual PIM multicast VPNs). PIM join load balancing is not supported on multiprotocol BGP-based multicast VPNs (also referred to as next-generation Layer 3 VPN multicast). When PIM join load balancing is enabled in a draft-rosen Layer 3 VPN scenario, the load balancing is achieved based on the join counts for the far-end PE routing devices, not for any intermediate P routing devices.

If an internal BGP (IBGP) multipath forwarding VPN route is available, the Junos OS uses the multipath forwarding VPN route to send join messages to the remote PE routers to achieve load balancing over the VPN.

By default, when multiple PIM joins are received for different groups, all joins are sent to the same upstream gateway chosen by the unicast routing protocol. Even if there are

multiple equal-cost paths available, these alternative paths are not utilized to distribute multicast traffic from the source to the various groups.

When PIM join load balancing is configured, the PIM joins are distributed equally among all equal-cost upstream interfaces and neighbors. Every new join triggers the selection of the least-loaded upstream interface and neighbor. If there are multiple neighbors on the same interface (for example, on a LAN), join load balancing maintains a value for each of the neighbors and distributes multicast joins (and downstream traffic) among these as well.

Join counts for interfaces and neighbors are maintained globally, not on a per-source basis. Therefore, there is no guarantee that joins for a particular source are load-balanced. However, the joins for all sources and all groups known to the routing device are load-balanced. There is also no way to administratively give preference to one neighbor over another: all equal-cost paths are treated the same way.

You can configure message filtering globally or for a routing instance. This example shows the global configuration.

You configure PIM join load balancing on the non-RP routers in the PIM domain.

1. Determine if there are multiple paths available for a source (for example, an RP) with the output of the **show pim join extensive** or **show pim source** commands.

```
user@host> show pim join extensive
Instance: PIM.master Family: INET

Group: 224.1.1.1
    Source: *
    RP: 10.255.245.6
    Flags: sparse,rptree,wildcard
    Upstream interface: t1-0/2/3.0
    Upstream neighbor: 192.168.38.57
    Upstream state: Join to RP
    Downstream neighbors:
        Interface: t1-0/2/1.0
            192.168.38.16 State: JOIN Flags; SRW Timeout: 164
Group: 224.2.127.254
    Source: *
    RP: 10.255.245.6
    Flags: sparse,rptree,wildcard
    Upstream interface: so-0/3/0.0
    Upstream neighbor: 192.168.38.47
    Upstream state: Join to RP
    Downstream neighbors:
        Interface: t1-0/2/3.0
            192.168.38.16 State: JOIN Flags; SRW Timeout: 164
```

   Note that for this router, the RP at IP address 10.255.245.6 is the source for two multicast groups: 224.1.1.1 and 224.2.127.254. This router has two equal-cost paths through two different upstream interfaces (**t1-0/2/3.0** and **so-0/3/0.0**) with two different neighbors (192.168.38.57 and 192.168.38.47). This router is a good candidate for PIM join load balancing.

2. On the non-RP router, configure PIM join load balancing.

   [edit protocols pim rp]

```
user@host# set static address 10.10.10.1
user@host# set interface all mode sparse version 2
user@host# set join-load-balance
```

The static address is the address of the RP.

3. Monitor the operation.

   If load balancing is enabled for this router, the number of PIM joins sent on each interface is shown in the output for the **show pim interfaces** command.

```
user@host> show pim interfaces
Instance: PIM.master

Name             Stat Mode        IP V State NbrCnt JoinCnt   DR address
lo0.0            Up   Sparse       4 2 DR         0       0   10.255.168.58
pe-1/2/0.32769   Up   Sparse       4 2 P2P        0       0
so-0/3/0.0       Up   Sparse       4 2 P2P        1       1
t1-0/2/1.0       Up   Sparse       4 2 P2P        1       0
t1-0/2/3.0       Up   Sparse       4 2 P2P        1       1
lo0.0            Up   Sparse       6 2 DR         0       0   fe80::2a0:a5ff:4b7
```

Note that the two equal-cost paths shown by the **show pim interfaces** command now have nonzero join counts. If the counts differ by more than one and were zero (0) when load balancing commenced, an error occurs (joins before load balancing are not redistributed). The join count also appears in the **show pim neighbors detail** output:

```
user@host> show pim neighbors detail
Interface: so-0/3/0.0

    Address: 192.168.38.46, IPv4, PIM v2, Mode: Sparse, Join Count: 0
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1689116164
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

    Address: 192.168.38.47, IPv4, PIM v2, Join Count: 1
        BFD: Disabled
        Hello Option Holdtime: 105 seconds 102 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 792890329
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Interface: t1-0/2/3.0

    Address: 192.168.38.56, IPv4, PIM v2, Mode: Sparse, Join Count: 0
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 678582286
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

    Address: 192.168.38.57, IPv4, PIM v2, Join Count: 1
        BFD: Disabled
        Hello Option Holdtime: 105 seconds 97 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1854475503
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Note that the join count is nonzero on the two load-balanced interfaces toward the upstream neighbors.

PIM join load balancing only takes effect when the feature is configured. Prior joins are not redistributed to achieve perfect load balancing. In addition, if an interface or neighbor fails, the new joins are redistributed among remaining active interfaces and neighbors. However, when the interface or neighbor is restored, prior joins are not redistributed. The **clear pim join-distribution** command redistributes the existing flows to new or restored upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you perform PIM join redistribution during a maintenance window.

Related Documentation

- *clear pim join-distribution* in the CLI Explorer
- show pim interfaces on page 417 in the CLI Explorer
- show pim neighbors on page 434 in the CLI Explorer
- show pim source on page 445 in the CLI Explorer

Modifying the Join State Timeout

This section describes how to configure the join state timeout.

A downstream router periodically sends join messages to refresh the join state on the upstream router. If the join state is not refreshed before the timeout expires, the join state is removed.

By default, the join state timeout is 210 seconds. You can change this timeout to allow additional time to receive the join messages. Because the messages are called join-prune messages, the name used is the **join-prune-timeout** statement.

To modify the timeout, include the **join-prune-timeout** statement:

```
user@host# set protocols pim join-prune-timeout 230
```

The join timeout value can be from 210 through 240 seconds.

Related Documentation

- join-prune-timeout on page 204

Example: Enabling Join Suppression

This example describes how to enable PIM join suppression.

Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See "Enabling PIM Sparse Mode" on page 65.

Overview

PIM join suppression enables a router on a multiaccess network to defer sending join messages to an upstream router when it sees identical join messages on the same network. Eventually, only one router sends these join messages, and the other routers suppress identical messages. Limiting the number of join messages improves scalability and efficiency by reducing the number of messages sent to the same router.

This example includes the following statements:

- **override-interval**—Sets the maximum time in milliseconds to delay sending override join messages. When a router sees a prune message for a join it is currently suppressing, it waits before it sends an override join message. Waiting helps avoid multiple downstream routers sending override join messages at the same time. The override interval is a random timer with a value of 0 through the maximum override value.
- **propagation-delay**—Sets a value in milliseconds for a prune pending timer, which specifies how long to wait before executing a prune on an upstream router. During this period, the router waits for any prune override join messages that might be currently suppressed. The period for the prune pending timer is the sum of the **override-interval** value and the value specified for **propagation-delay**.
- **reset-tracking-bit**—Enables PIM join suppression on each multiaccess downstream interface. This statement resets a tracking bit field (T-bit) on the LAN prune delay hello option from the default of 1 (join suppression disabled) to 0 (join suppression enabled).

  When multiple identical join messages are received, a random join suppression timer is activated, with a range of 66 through 84 milliseconds. The timer is reset each time join suppression is triggered.

Figure 14 on page 71 shows the topology used in this example.

Figure 14: Join Suppression

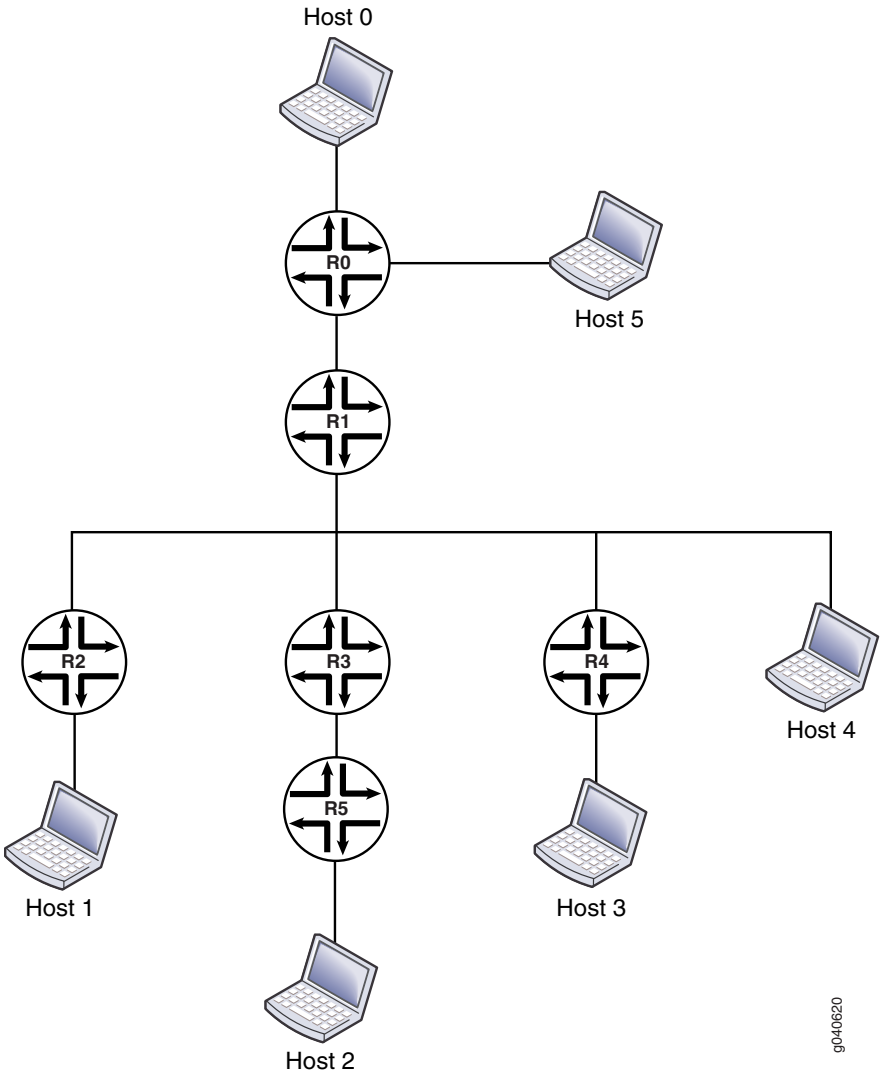

The items in the figure represent the following functions:

- Host 0 is the multicast source.
- Host 1, Host 2, Host 3, and Host 4 are receivers.
- Router R0 is the first-hop router and the RP.
- Router R1 is an upstream router.
- Routers R2, R3, R4, and R5 are downstream routers in the multicast LAN.

This example shows the configuration of the downstream devices: Routers R2, R3, R4, and R5.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
[edit]
set protocols pim traceoptions file pim.log
set protocols pim traceoptions file size 5m
set protocols pim traceoptions file world-readable
set protocols pim traceoptions flag join detail
set protocols pim traceoptions flag prune detail
set protocols pim traceoptions flag normal detail
set protocols pim traceoptions flag register detail
set protocols pim rp static address 10.255.112.160
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
set protocols pim reset-tracking-bit
set protocols pim propagation-delay 500
set protocols pim override-interval 4000
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure PIM join suppression on a non-RP downstream router in the multicast LAN:

1. Configure PIM sparse mode on the interfaces.

```
[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set rp static address 10.255.112.160
[edit protocols pim]
user@host# set interface all mode sparse version 2
[edit protocols pim]
user@host# set interface all version 2
[edit protocols pim]
user@host# set interface fxp0.0 disable
```

2. Enable the join suppression timer.

```
[edit protocols pim]
user@host# set reset-tracking-bit
```

3. Configure the prune override interval value.

```
[edit protocols pim]
user@host# set override-interval 4000
```

4. Configure the propagation delay of the link.

```
[edit protocols pim]
user@host# set propagation-delay 500
```

5. (Optional) Configure PIM tracing operations.

```
[edit protocols pim]
user@host# set traceoptions file pim.log size 5m world-readable
[edit protocols pim]
user@host# set traceoptions flag join detail
[edit protocols pim]
user@host# set traceoptions flag normal detail
[edit protocols pim]
user@host# set traceoptions flag register detail
```

6. If you are done configuring the device, commit the configuration.

```
[edit protocols pim]
user@host# commit
```

Results

From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols
pim {
    traceoptions {
        file pim.log size 5m world-readable;
        flag join detail;
        flag prune detail;
        flag normal detail;
        flag register detail;
    }
    rp {
        static {
            address 10.255.112.160;
        }
    }
    interface all {
        mode sparse;
        version 2;
    }
    interface fxp0.0 {
        disable;
    }
    reset-tracking-bit;
    propagation-delay 500;
    override-interval 4000;
}
```

Verification

To verify the configuration, run the following commands on the upstream and downstream routers:

- show pim join **extensive**

- **show multicast route extensive**

Related Documentation

- Example: Configuring the PIM Assert Timeout on page 99
- *Example: Configuring PIM RPF Selection*
- Example: Configuring the PIM SPT Threshold Policy on page 101
- Enabling PIM Sparse Mode on page 65
- PIM Overview on page 3

CHAPTER 17

Static RP

Configuring Local PIM RPs

Local RP configuration makes the routing device a statically defined RP. Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You can configure a local RP globally or for a routing instance. This example shows how to configure a local RP in a routing instance for IPv4 or IPv6.

To configure the routing device's RP properties:

1. Configure the routing instance as the local RP.

   ```
   [routing-instances VPN-A protocols pim]
   user@host# set rp local
   ```

2. Configure the IP protocol family and IP address.

   IPv6 PIM hello messages are sent to every interface on which you configure **family inet6**, whether at the PIM level of the hierarchy or not. As a result, if you configure an interface with both **family inet** at the **[edit interface *interface-name*]** hierarchy level and **family inet6** at the **[edit protocols pim interface *interface-name*]** hierarchy level, PIM sends both IPv4 and IPv6 hellos to that interface.

   By default, PIM operates in sparse mode on an interface. If you explicitly configure sparse mode, PIM uses this setting for all IPv6 multicast groups. However, if you configure sparse-dense mode, PIM does not accept IPv6 multicast groups as dense groups and operates in sparse mode over them.

   ```
   [edit routing-instances VPN-A protocols pim rp local]
   user@host# set family inet6 address 2001:db8:85a3::8a2e:370:7334
   user@host# set family inet address 10.1.2.254
   ```

3. (IPv4 only) Configure the routing device's RP priority.

NOTE: The priority statement is not supported for IPv6, but is included here for informational purposes. The routing device's priority value for becoming the RP is included in the bootstrap messages that the routing device sends. Use a smaller number to increase the likelihood that the routing device becomes the RP for local multicast groups. Each PIM routing device uses the priority value and other factors to determine the candidate RPs for a particular group range. After the set of candidate RPs is distributed, each routing device determines algorithmically the RP from the candidate RP set using a hash function. By default, the priority value is set to 1. If this value is set to 0, the bootstrap router can override the group range being advertised by the candidate RP.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set priority 5
```

4. Configure the groups for which the routing device is the RP.

   By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which this routing device can be the RP.

   ```
   [edit routing-instances VPN-A protocols pim rp local]
   user@host# set group-ranges fec0::/10
   user@host# set group-ranges 10.1.2.0/24
   ```

5. (IPv4 only) Modify the local RP hold time.

   If the local routing device is configured as an RP, it is considered a candidate RP for its local multicast groups. For candidate RPs, the hold time is used by the bootstrap router to time out RPs, and applies to the bootstrap RP-set mechanism. The RP hold time is part of the candidate RP advertisement message sent by the local routing device to the bootstrap router. If the bootstrap router does not receive a candidate RP advertisement from an RP within the hold time, it removes that routing device from its list of candidate RPs. The default hold time is 150 seconds.

   ```
   [edit routing-instances VPN-A protocols pim rp local]
   user@host# set hold-time 200
   ```

6. (Optional) Override dynamic RP for the specified group address range.

   If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

   If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

   ```
   [edit routing-instances VPN-A protocols pim rp local]
   user@host# set override
   ```

7. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

Related Documentation

- PIM Overview on page 3
- *Understanding MLD*

Configuring the Static PIM RP Address on the Non-RP Routing Device

Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You configure a static RP address on the non-RP routing device. This enables the non-RP routing device to recognize the local statically defined RP. For example, if R0 is a non-RP router and R1 is the local RP router, you configure R0 with the static RP address of R1. The static IP address is the routable address assigned to the loopback interface on R1. In the following example, the loopback address of the RP is 2001:db8:85a3::8a2e:370:7334.

You can configure a static RP address globally or for a routing instance. This example shows how to configure a static RP address in a routing instance for IPv6.

To configure the static RP address:

1. On a non-RP routing device, configure the routing instance to point to the routable address assigned to the loopback interface of the RP.

   ```
   [routing-instances VPN-A protocols pim rp]
   user@host# set static address 2001:db8:85a3::8a2e:370:7334
   ```

   NOTE: Logical systems are also supported. You can configure a static RP address in a logical system only if the logical system is not directly connected to a source.

2. (Optional) Set the PIM sparse mode version.

   For each static RP address, you can optionally specify the PIM version. The default PIM version is version 1.

   ```
   [edit routing-instances VPN-A protocols pim rp]
   user@host# set static address 2001:db8:85a3::8a2e:370:7334 version 2
   ```

   NOTE: The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIM version 1 is the default for RP mode ([edit pim rp static address *address*]). PIM version 2 is the default for interface mode ([edit pim interface *interface-name*]). Explicitly configured versions override the defaults.

3. (Optional) Set the group address range.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which the 2001:db8:85a3::8a2e:370:7334 address can be the RP.

```
[edit routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334 group-ranges fec0::/10
```

The RP that you select for a particular group must be consistent across all routers in a multicast domain.

4. (Optional) Override dynamic RP for the specified group address range.

   If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

   If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

   ```
   [edit routing-instances VPN-A protocols pim rp static address
     2001:db8:85a3::8a2e:370:7334]
   user@host# set override
   ```

5. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

Related Documentation

- PIM Overview on page 3
- *Understanding MLD*

CHAPTER 18

Anycast RP

Example: Configuring PIM Anycast With or Without MSDP

When you configure anycast RP, you bypass the restriction of having one active rendezvous point (RP) per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP stops operating, sources and receivers are taken to a new RP by means of unicast routing.

You can configure anycast RP to use PIM and MSDP for IPv4, or PIM alone for both IPv4 and IPv6 scenarios. Both are discussed in this section.

We recommend a static RP mapping with anycast RP over a bootstrap router and auto-RP configuration because it provides all the benefits of a bootstrap router and auto-RP without the complexity of the BSR and auto-RP mechanisms.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default RP mode (at the **[edit protocols pim rp static address *address*]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface *interface-name*]** hierarchy level). Explicitly configured versions override the defaults. This example explicitly configures PIMv2 on the interfaces.

The following example shows an anycast RP configuration for the RP routers, first with MSDP and then using PIM alone, and for non-RP routers.

1. For a network using an RP with MSDP, configure the RP using the **lo0** loopback interface, which is always up. Include the **address** statement and specify the unique and routable router ID and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement for the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32;
        primary;
        address 198.58.3.253/32;
      }
    }
  }
}
```

2. Specify the RP address. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level (the same address as the secondary **lo0** interface).

   For all interfaces, include the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
  pim {
    rp {
      local {
        family inet;
        address 198.58.3.253;
      }
      interface all {
        mode sparse;
        version 2;
      }
      interface fxp0.0 {
        disable;
      }
    }
  }
}
```

3. Configure MSDP peering. Include the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, include the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
  msdp {
    peer 198.58.3.250 {
```



```
        local-address address 198.58.3.254;
      }
    }
  }
```

NOTE: If you need to configure a PIM RP for both IPv4 and IPv6 scenarios, perform Step 4 and Step 5. Otherwise, go to Step 6.

4. Configure an RP using the **lo0** loopback interface, which is always up. Include the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement on the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32 {
          primary;
        }
        address 198.58.3.253/32;
      }
    }
  }
}
```

5. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

   For all interfaces, include the **mode** statement to set the mode to **sparse**, and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by Including the **disable** statement for that interface.

   Include the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

   MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

```
protocols {
  pim {
    rp {
      local {
        family inet {
          address 198.58.3.253;
          anycast-pim {
            rp-set {
```

```
                address 198.58.3.240;
                address 198.58.3.241 forward-msdp-sa;
              }
              local-address 198.58.3.254; #If not configured, use lo0 primary
            }
          }
        }
      }
      interface all {
        mode sparse;
        version 2;
      }
      interface fxp0.0 {
        disable;
      }
    }
  }
```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at the **[edit protocols pim rp static]** hierarchy level. Include the **version** statement at the **[edit protocols pim rp static address]** hierarchy level to specify PIM version 2.

```
protocols {
  pim {
    rp {
      static {
        address 198.58.3.253 {
          version 2;
        }
      }
    }
  }
}
```

7. Include the **mode** statement at the **[edit protocols pim interface all]** hierarchy level to specify sparse mode on all interfaces. Then include the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

Configuring a PIM Anycast RP Router with MSDP

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
protocols {
    pim {
        rp {
            local {
                family inet;
                address 198.58.3.253;
            }
            interface all {
                mode sparse;
                version 2;
            }
            interface fxp0.0 {
                disable;
            }
        }
    }
}
```

To configure MSDP peering, add the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, add the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
    msdp {
        peer 198.58.3.250 {
            local-address 198.58.3.254;
        }
    }
}
```

Configuring a PIM Anycast RP Router Using Only PIM

In this example, configure an RP using the **lo0** loopback interface, which is always up. Use the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this case, the router ID is 198.58.3.254/32 and the shared RP address is 198.58.3.253/32. Add the flag statement **primary** to the first address. Using this flag selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
    lo0 {
```

```
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32 {
          primary;
        }
        address 198.58.3.253/32;
      }
    }
  }
}
```

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse**, and include the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

Use the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

```
protocols {
  pim {
    rp {
      local {
        family inet {
          address 198.58.3.253;
          anycast-pim {
            rp-set {
              address 198.58.3.240;
              address 198.58.3.241 forward-msdp-sa;
            }
            local-address 198.58.3.254; #If not configured, use lo0 primary
          }
        }
      }
    }
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

Configuring All PIM Anycast Non-RP Routers

Use the **mode** statement at the **[edit protocols pim rp interface all]** hierarchy level to specify sparse mode on all interfaces. Then add the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

Example: Configuring Multiple RPs in a Domain with Anycast RP

This example shows how to configure anycast RP on each RP router in the PIM-SM domain. With this configuration you can deploy more than one RP for a single group range. This enables load balancing and redundancy.

Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See "Enabling PIM Sparse Mode" on page 65.

Overview

When you configure anycast RP, the RP routers in the PIM-SM domain use a shared address. In this example, the shared address is 10.1.1.2/32. Anycast RP uses Multicast Source Discovery Protocol (MSDP) to discover and maintain a consistent view of the active sources. Anycast RP also requires an RP selection method, such as static, auto-RP, or bootstrap RP. This example uses static RP and shows only one RP router configuration.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

RP Routers

```
set interfaces lo0 unit 0 family inet address 192.168.132.1/32 primary
set interfaces lo0 unit 0 family inet address 10.1.1.2/32
set protocols msdp local-address 192.168.132.1
set protocols msdp peer 192.168.12.1
set protocols pim rp local address 10.1.1.2
set routing-options router-id 192.168.132.1
```

Non-RP Routers

```
set protocols pim rp static address 10.1.1.2
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure anycast RP:

1. On each RP router in the domain, configure the shared anycast address on the router's loopback address.

   ```
   [edit interfaces]
   user@host# set lo0 unit 0 family inet address 10.1.1.2/32
   ```

2. On each RP router in the domain, make sure that the router's regular loopback address is the primary address for the interface, and set the router ID.

   ```
   [edit interfaces]
   user@host# set lo0 unit 0 family inet address 192.168.132.1/32 primary

   [edit routing-options]
   user@host# set router-id 192.168.132.1
   ```

3. On each RP router in the domain, configure the local RP address, using the shared address.

   ```
   [edit protocols pim]
   user@host# set rp local address 10.1.1.2
   ```

4. On each RP router in the domain, create MSDP sessions to the other RPs in the domain.

   ```
   [edit protocols msdp]
   user@host# set local-address 192.168.132.1
   user@host# set peer 192.168.12.1
   ```

5. On each non-RP router in the domain, configure a static RP address using the shared address.

   ```
   [edit protocols pim]
   user@host# set rp static address 10.1.1.2
   ```

6. If you are done configuring the devices, commit the configuration.

```
user@host# commit
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
lo0 {
    unit 0 {
        family inet {
            address 192.168.132.1/32 {
                primary;
            }
            address 10.1.1.2/32;
        }
    }
}
```

On the RP routers:

```
user@host# show protocols
msdp {
    local-address 192.168.132.1;
    peer 192.168.12.1;
}
pim {
    rp {
        local {
            address 10.1.1.2;
        }
    }
}
```

On the non-RP routers:

```
user@host# show protocols
pim {
    rp {
        static {
            address 10.1.1.2;
        }
    }
}
```

```
user@host# show routing-options
router-id 192.168.132.1;
```

Verification

To verify the configuration, run the show pim rps **extensive inet** command.

Related Documentation

- Example: Configuring PIM Anycast With or Without MSDP on page 79
- Understanding PIM Sparse Mode on page 7
- Understanding RP Mapping with Anycast RP on page 13

CHAPTER 19

PIM Bootstrap Router

Configuring PIM Bootstrap Properties for IPv4 or IPv6

For correct operation, every multicast router within a PIM domain must be able to map a particular multicast group address to the same rendezvous point (RP). The bootstrap router mechanism is one way that a multicast router can learn the set of group-to-RP mappings. Bootstrap routers are supported in IPv4 and IPv6.

To determine which routing device is the RP, all routing devices within a PIM domain collect bootstrap messages. A PIM domain is a contiguous set of routing devices that implement PIM. All devices are configured to operate within a common boundary. The domain's bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

You can configure bootstrap properties globally or for a routing instance. This example shows the global configuration.

To configure the bootstrap router properties:

1. Configure the bootstrap priority.

   By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap router. The routing device with the highest priority value is elected to be the bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router. A simple bootstrap configuration assigns a bootstrap priority value to a routing device.

NOTE: In the IPv4-only configuration, specifying a bootstrap priority of 0 disables the bootstrap function and does not cause the routing device to send BSR packets with a 0 in the priority field. In the combined IPv4 and IPv6 configuration, specifying a bootstrap priority of 0 does not disable the function, but causes the routing device to send BSR packets with a 0 in the priority field. To disable the bootstrap function in the IPv4 and IPv6 configuration, delete the **bootstrap** statement.

```
user@host# edit protocols pim rp
user@host# set bootstrap family inet priority 3
```

2. (Optional) Create import and export policies to control the flow of bootstrap messages to and from the RP, and apply the policies to PIM. Import and export policies are useful when some of the routers in your PIM domain have interfaces that connect to other PIM domains. Configuring a policy prevents bootstrap messages from crossing domain boundaries. The **import** statement prevents messages from being imported into the RP. The **export** statement prevents messages from being exported from the RP.

```
[edit protocols pim rp]
user@host# set bootstrap family inet import pim-bootstrap-import
user@host# set bootstrap family inet export pim-bootstrap-export
user@host# exit
```

3. Configure the policies.

```
user@host# edit policy-options policy-statement pim-bootstrap-import
[edit policy-options policy-statement pim-bootstrap-import]
user@host# set from interface se-0/0/0
user@host# set then reject
user@host# exit
user@host# edit policy-options policy-statement pim-bootstrap-export
user@host# set from interface se-0/0/0
user@host# set then reject
user@host# exit
```

4. Monitor the operation of PIM bootstrap routers by running the **show pim bootstrap** command.

Related Documentation

- Understanding PIM Sparse Mode on page 7
- Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 90
- **show pim bootstrap** on page 415 in the CLI Explorer

Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain

In this example, the **from interface so-0-1/0 then reject** policy statement rejects bootstrap messages from the specified interface (the example is configured for both IPv4 and IPv6 operation):

```
protocols {
```

```
    pim {
      rp {
        bootstrap {
          family inet {
            priority 1;
            import pim-import;
            export pim-export;
          }
          family inet6 {
            priority 1;
            import pim-import;
            export pim-export;
          }
        }
      }
    }
}
policy-options {
    policy-statement pim-import {
      from interface so-0/1/0;
      then reject;
    }
    policy-statement pim-export {
      to interface so-0/1/0;
      then reject;
    }
}
```

Example: Configuring PIM BSR Filters

Configure a filter to prevent BSR messages from entering or leaving your network. Add this configuration to all routers:

```
protocols {
    pim {
      rp {
        bootstrap-import no-bsr;
        bootstrap-export no-bsr;
      }
    }
}
policy-options {
    policy-statement no-bsr {
      then reject;
    }
}
```

CHAPTER 20

PIM Filtering

Configuring Interface-Level PIM Neighbor Policies

You can configure a policy to filter unwanted PIM neighbors. In the following example, the PIM interface compares neighbor IP addresses with the IP address in the policy statement before any hello processing takes place. If any of the neighbor IP addresses (primary or secondary) match the IP address specified in the prefix list, PIM drops the hello packet and rejects the neighbor.

If you configure a PIM neighbor policy after PIM has already established a neighbor adjacency to an unwanted PIM neighbor, the adjacency remains intact until the neighbor hold time expires. When the unwanted neighbor sends another hello message to update its adjacency, the router recognizes the unwanted address and rejects the neighbor.

To configure a policy to filter unwanted PIM neighbors:

1. Configure the policy. The neighbor policy must be a properly structured policy statement that uses a prefix list (or a route filter) containing the neighbor primary address (or any secondary IP addresses) in a prefix list, and the **reject** option to reject the unwanted address.

   ```
   [edit policy-options]
   user@host# set prefix-list nbrGroup 1 20.20.20.1/32
   user@host# set policy-statement nbr-policy from prefix-list nbrGroup1
   user@host# set policy-statement nbr-policy then reject
   ```

2. Configure the interface globally or in the routing instance. This example shows the configuration for the routing instance.

   ```
   [edit routing-instances PIM.master protocols pim]
   user@host# set neighbor-policy nbr-policy
   ```

3. Verify the configuration by checking the **Hello dropped on neighbor policy** field in the output of the **show pim statistics** command.

Related Documentation

- Understanding PIM Sparse Mode on page 7
- *Routing Policy Feature Guide for Routing Devices*
- **show pim statistics on page 448** in the CLI Explorer

Filtering Outgoing PIM Join Messages

When the core of your network is using MPLS, PIM join and prune messages stop at the customer edge (CE) routers and are not forwarded toward the core, because these routers do not have PIM neighbors on the core-facing interfaces. When the core of your network is using IP, PIM join and prune messages are forwarded to the upstream PIM neighbors in the core of the network.

When the core of your network is using a mix of IP and MPLS, you might want to filter certain PIM join and prune messages at the upstream egress interface of the CE routers.

You can filter PIM sparse mode (PIM-SM) join and prune messages at the egress interfaces for IPv4 and IPv6 in the upstream direction. The messages can be filtered based on the group address, source address, outgoing interface, PIM neighbor, or a combination of these values. If the filter is removed, the join is sent after the PIM periodic join timer expires.

To filter PIM sparse mode join and prune messages at the egress interfaces, create a policy rejecting the group address, source address, outgoing interface, or PIM neighbor, and then apply the policy.

The following example filters PIM join and prune messages for group addresses 224.0.1.2 and 225.1.1.1.

1. In configuration mode, create the policy.

```
user@host# set policy-options policy-statement block-groups term t1 from route-filter
  224.0.1.2/32 exact
user@host# set policy-options policy-statement block-groups term t1 from route-filter
  225.1.1.1/32 exact
user@host# set policy-options policy-statement block-groups term t1 then reject
user@host# set policy-options policy-statement block-groups term last then accept
```

2. Verify the policy configuration by running the **show policy-options** command.

```
user@host# show policy-options
policy-statement block-groups {
  term t1 {
    from {
      route-filter 224.0.1.2/32 exact;
      route-filter 225.1.1.1/32 exact;
      then reject;
    }
    term last {
      then accept;
    }
  }
```

3. Apply the PIM join and prune message filter.

```
user@host> set protocols pim export block-groups
```

4. After the configuration is committed, use the **show pim statistics** command to verify that outgoing PIM join and prune messages are being filtered.

```
user@host> show pim statistics | grep filtered
RP Filtered Source                  0

Rx Joins/Prunes filtered            0

Tx Joins/Prunes filtered            254
```

The egress filter count is shown on the **Tx Joins/Prunes filtered** line.

Related Documentation

- Filtering Incoming PIM Join Messages on page 95

Filtering Incoming PIM Join Messages

Multicast scoping controls the propagation of multicast messages. Whereas multicast scoping prevents the actual multicast data packets from flowing in or out of an interface, PIM join filters prevent a state from being created in a router. A state—the (*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. Using PIM join filters prevents the transport of multicast traffic across a network and the dropping of packets at a scope at the edge of the network. Also, PIM join filters reduce the potential for denial-of-service (DoS) attacks and PIM state explosion—large numbers of PIM join messages forwarded to each router on the rendezvous-point tree (RPT), resulting in memory consumption.

To use PIM join filters to efficiently restrict multicast traffic from certain source addresses, create and apply the routing policy across all routers in the network.

See Table 4 on page 95 for a list of match conditions.

Table 4: PIM Join Filter Match Conditions

Match Condition	Matches On
interface	Router interface or interfaces specified by name or IP address
neighbor	Neighbor address (the source address in the IP header of the join and prune message)
route-filter	Multicast group address embedded in the join and prune message
source-address-filter	Multicast source address embedded in the join and prune message

The following example shows how to create a PIM join filter. The filter is composed of a route filter and a source address filter—**bad-groups** and **bad-sources**, respectively. the **bad-groups** filter prevents (*,G) or (S,G) join messages from being received for all groups listed. The **bad-sources** filter prevents (S,G) join messages from being received for all sources listed. The **bad-groups** filter and **bad-sources** filter are in two different terms. If route filters and source address filters are in the same term, they are logically ANDed.

To filter incoming PIM join messages:

1. Configure the policy.

   ```
   [edit policy-statement pim-join-filter term bad-groups]
   user@host# set from route-filter 224.0.1.2/32 exact
   user@host# set from route-filter 239.0.0.0/8 orlonger
   user@host# set then reject

   [edit policy-statement pim-join-filter term bad-sources]
   user@host# set from source-address-filter 10.0.0.0/8 orlonger
   user@host# set from source-address-filter 127.0.0.0/8 orlonger
   user@host# set then reject

   [edit policy-statement pim-join-filter term last]
   user@host# set then accept
   ```

2. Apply one or more policies to routes being imported into the routing table from PIM.

   ```
   [edit protocols pim]
   user@host# set import pim-join-filter
   ```

3. Verify the configuration by checking the output of the **show pim join** and **show policy** commands.

Related Documentation

- *Understanding Multicast Administrative Scoping*
- Filtering Outgoing PIM Join Messages on page 94
- **show pim join on page 420** in the CLI Explorer
- *show policy* in the CLI Explorer

Configuring Register Message Filters on a PIM RP and DR

PIM register messages are sent to the rendezvous point (RP) by a designated router (DR). When a source for a group starts transmitting, the DR sends unicast PIM register packets to the RP.

Register messages have the following purposes:

- Notify the RP that a source is sending to a group.
- Deliver the initial multicast packets sent by the source to the RP for delivery down the shortest-path tree (SPT).

The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, you want more control over which sources an RP discovers, or which sources a DR notifies other RPs about. A high degree of control over PIM register messages is provided by RP or DR register message filtering. Message filtering prevents unauthorized groups and sources from registering with an RP router.

You configure RP or DR register message filtering to control the number and location of multicast sources that an RP discovers. You can apply register message filters on a DR to control outgoing register messages, or apply them on an RP to control incoming register messages.

When anycast RP is configured, all RPs in the anycast RP set need to be configured with the same register message filtering policy.

You can configure message filtering globally or for a routing instance. These examples show the global configuration.

To configure an RP filter to drop the register packets for multicast group range 224.1.1.0/24 from source address 10.10.94.2:

1. On the RP, configure the policy.

```
[edit policy-options policy-statement incoming-policy-for-rp from]
user@host# set route-filter 224.1.1.0/24 orlonger
user@host# set source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy incoming-policy-for-rp
user@host# set local address 10.10.10.5
user@host# exit
```

To configure a DR filter to prevent sending register packets for group range 224.1.1.0/24 and source address 10.10.10.1/32:

1. On the DR, configure the policy.

```
[edit policy-options policy-statement outgoing-policy-for-rp]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.10.1/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the DR.

   The static address is the address of the RP to which you do not want the DR to send the filtered register messages.

```
[edit protocols pim rp]
user@host# set dr-register-policy outgoing-policy-for-dr
user@host# set static 10.10.10.3
user@host# exit
```

To configure a policy expression to accept register messages for multicast group 224.1.1.5 but reject those for 224.1.1.1:

1. On the RP, configure the policies.

```
[edit policy-options policy-statement reject_224_1_1_1]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit

[edit policy-options policy-statement accept_224_1_1_5]
user@host# set term one from route-filter 224.1.1.5/32 exact
```

```
user@host# set term one from source-address-filter 10.10.94.2/32 exact
user@host# set term one then accept
user@host# set term two then reject
user@host# exit
```

2. Apply the policies to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy [ reject_224_1_1_1 | accept_224_1_1_5 ]
user@host# set local address 10.10.10.5
```

To monitor the operation of the filters, run the **show pim statistics** command. The command output contains the following fields related to filtering:

- **RP Filtered Source**
- **Rx Joins/Prunes filtered**
- **Tx Joins/Prunes filtered**
- **Rx Register msgs filtering drop**
- **Tx Register msgs filtering drop**

Related Documentation

- PIM Sparse Mode Source Registration on page 23
- Filtering RP and DR Register Messages on page 18
- **show pim statistics on page 448** in the CLI Explorer

CHAPTER 21

PIM RPT and SPT Cutover

Example: Configuring the PIM Assert Timeout

This example shows how to configure the timeout period for a PIM assert forwarder.

Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See "Enabling PIM Sparse Mode" on page 65.

Overview

The role of PIM assert messages is to determine the forwarder on a network with multiple routers. The forwarder is the router that forwards multicast packets to a network with multicast group members. The forwarder is generally the same as the PIM DR.

A router sends an assert message when it receives a multicast packet on an interface that is listed in the outgoing interface list of the matching routing entry. Receiving a message on an outgoing interface is an indication that more than one router forwards the same multicast packets to a network.

In Figure 15 on page 100, both routing devices R1 and R2 forward multicast packets for the same (S,G) entry on a network. Both devices detect this situation and both devices send assert messages on the Ethernet network. An assert message contains, in addition to a source address and group address, a unicast cost metric for sending packets to the

source, and a preference metric for the unicast cost. The preference metric expresses a preference between unicast routing protocols. The routing device with the smallest preference metric becomes the forwarder (also called the assert winner). If the preference metrics are equal, the device that sent the lowest unicast cost metric becomes the forwarder. If the unicast metrics are also equal, the routing device with the highest IP address becomes the forwarder. After the transmission of assert messages, only the forwarder continues to forward messages on the network.

When an assert message is received and the RPF neighbor is changed to the assert winner, the assert timer is set to an assert timeout period. The assert timeout period is restarted every time a subsequent assert message for the route entry is received on the incoming interface. When the assert timer expires, the routing device resets its RPF neighbor according to its unicast routing table. Then, if multiple forwarders still exist, the forwarders reenter the assert message cycle. In effect, the assert timeout period determines how often multicast routing devices enter a PIM assert message cycle.

The range is from 5 through 210 seconds. The default is 180 seconds.

Assert messages are useful for LANs that connect multiple routing devices and no hosts.

Figure 15 on page 100 shows the topology for this example.

Figure 15: PIM Assert Topology

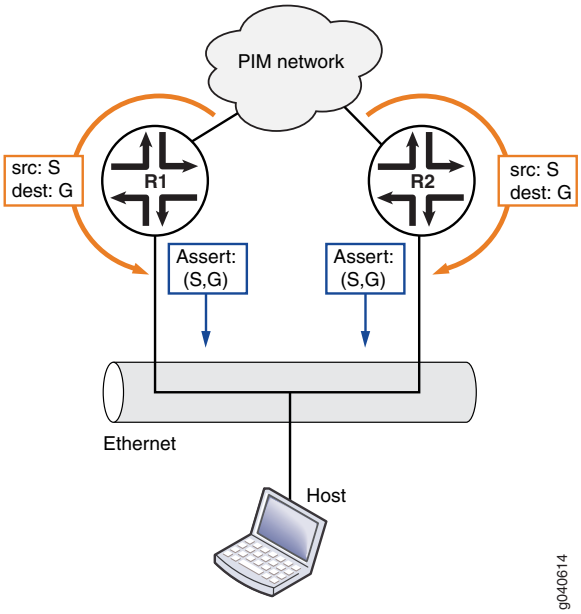

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an assert timeout:

1. Configure the timeout period, in seconds.

   ```
   [edit protocols pim]
   user@host# set assert-timeout 60
   ```

2. (Optional) Trace assert messages.

   ```
   [edit protocols pim]
   user@host# set traceoptions file PIM.log
   user@host# set traceoptions flag assert detail
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   user@host# commit
   ```

4. To verify the configuration, run the following commands:

   - show pim join
   - show pim statistics

Related Documentation

- Configuring PIM Trace Options on page 57
- SPT Cutover on page 27
- SPT Cutover Control on page 30

Example: Configuring the PIM SPT Threshold Policy

This example shows how to apply a policy that suppresses the transition from the rendezvous-point tree (RPT) rooted at the RP to the shortest-path tree (SPT) rooted at the source.

Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.

- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See "Enabling PIM Sparse Mode" on page 65.

Overview

Multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or through an SPT rooted at the source. In some cases, the last-hop routing device needs to stay on the shared RPT to the RP and not transition to a direct SPT to the source. Receiving the multicast data traffic on SPT is optimal but introduces more state in the network, which might not be desirable in some multicast deployments. Ideally, low-bandwidth multicast streams can be forwarded on the SPT, and high-bandwidth streams can use the SPT. This example shows how to configure such a policy.

This example includes the following settings:

- **spt-threshold**—Enables you to configure an SPT threshold policy on the last-hop routing device to control the transition to a direct SPT. When you include this statement in the main PIM instance, the PE router stays on the RPT for control traffic.
- **infinity**—Applies an SPT cutover threshold of infinity to a source-group address pair, so that the last-hop routing device never transitions to a direct SPT. For all other source-group address pairs, the last-hop routing device transitions immediately to a direct SPT rooted at the source DR. This statement must reference a properly configured policy to set the SPT cutover threshold for a particular source-group pair to infinity. The use of values other than infinity for the SPT threshold is not supported. You can configure more than one policy.
- **policy-statement**—Configures the policy. The simplest type of SPT threshold policy uses a route filter and source address filter to specify the multicast group and source addresses and to set the SPT threshold for that pair of addresses to infinity. The policy is applied to the main PIM instance.

  This example sets the SPT transition value for the source-group pair 10.10.10.1 and 224.1.1.1 to infinity. When the policy is applied to the last-hop router, multicast traffic from this source-group pair never transitions to a direct SPT to the source. Traffic will continue to arrive through the RP. However, traffic for any other source-group address combination at this router transitions to a direct SPT to the source.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.
- When the policy is configured for the first time, the routing device continues to transition to the direct SPT for the source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- If you do not clear the PIM-join state when you apply the infinity policy configuration for the first time, you must apply it before the PE router is brought up.
- When the policy is deleted for a source-group address pair for the first time, the routing device does not transition to the direct SPT for that source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- When the policy is changed for a source-group address pair for the first time, the routing device does not use the new policy until the PIM-join state is cleared with the **clear pim join** command.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
[edit]
set policy-options policy-statement spt-infinity-policy term one from route-filter
   224.1.1.1/32 exact
set policy-options policy-statement spt-infinity-policy term one from source-address-filter
   10.10.10.1/32 exact
set policy-options policy-statement spt-infinity-policy term one then accept
set policy-options policy-statement spt-infinity-policy term two then reject
set protocols pim spt-threshold infinity spt-infinity-policy
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure an SPT threshold policy:

1. Apply the policy.

```
[edit]
user@host# edit protocols pim
```

```
[edit protocols pim]
user@host# set spt-threshold infinity spt-infinity-policy
[edit protocols pim]
user@host# exit
```

2. Configure the policy.

```
[edit]
user@host# edit policy-options policy-statement spt-infinity-policy
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from route-filter 224.1.1.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from source-address-filter 10.10.10.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one then accept
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term two then reject
[edit policy-options policy-statement spt-infinity-policy]
user@host# exit
policy-statement {
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

4. Clear the PIM join cache to force the configuration to take effect.

```
[edit]
user@host# run clear pim join
```

Results

Confirm your configuration by entering the **show policy-options** command and the **show protocols** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement spt-infinity-policy {
  term one {
    from {
      route-filter 224.1.1.1/32 exact;
      source-address-filter 10.10.10.1/32 exact;
    }
    then accept;
  }
  term two {
    then reject;
  }
}
```

```
user@host# show protocols
pim {
  spt-threshold {
    infinity spt-infinity-policy;
  }
}
```

Verification

To verify the configuration, run the **show pim join** command.

Related Documentation

- SPT Cutover Control on page 30

CHAPTER 22

PIM and the BFD Protocol

Configuring BFD for PIM

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the Protocol Independent Multicast (PIM) hello hold time, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

You must specify the minimum transmit and minimum receive intervals to enable BFD on PIM.

To enable failure detection:

1. Configure the interface globally or in a routing instance.

   This example shows the global configuration.

   ```
   [edit protocols pim]
   user@host# edit interface fe-1/0/0.0 family inet bfd-liveness-detection
   ```

2. Configure the minimum transmit interval.

This is the minimum interval after which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set transmit-interval 350
```

3. Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session.

   Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

   ```
   [edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
   user@host# set minimum-receive-interval 350
   ```

4. (Optional) Configure other BFD settings.

   As an alternative to setting the receive and transmit intervals separately, configure one interval for both.

   ```
   [edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
   user@host# set minimum-interval 350
   ```

5. Configure the threshold for the adaptation of the BFD session detection time.

   When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

   ```
   [edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
   user@host# set detection-time threshold 800
   ```

6. Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

   ```
   [edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
   user@host# set multiplier 50
   ```

7. Configure the BFD version.

   ```
   [edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
   user@host# set version 1
   ```

8. Specify that BFD sessions should not adapt to changing network conditions.

   We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

   ```
   [edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
   user@host# set no-adaptation
   ```

9. Verify the configuration by checking the output of the **show bfd session** command.

Related Documentation

- *show bfd session* in the CLI Explorer

Configuring BFD Authentication for PIM

Beginning with Junos OS Release 9.6, you can configure authentication for Bidirectional Forwarding Detection (BFD) sessions running over Protocol Independent Multicast (PIM).

Routing instances are also supported. The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the PIM protocol.
2. Associate the authentication keychain with the PIM protocol.
3. Configure the related security authentication keychain.

The following sections provide instructions for configuring and viewing BFD authentication on PIM:

- Configuring BFD Authentication Parameters on page 109
- Viewing Authentication Information for BFD Sessions on page 110

Configuring BFD Authentication Parameters

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on a PIM route or routing instance.

   ```
   [edit protocols pim]
   user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication
     algorithm keyed-sha-1
   ```

   NOTE: Nonstop active routing (NSR) is not supported with the meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on the specified PIM route or routing instance with the unique security authentication keychain attributes.

   The keychain you specify must match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

   ```
   [edit protocols pim]
   user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication
     keychain bfd-pim
   ```

   NOTE: The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 2.
- At least one key, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
user@host# set authentication-key-chains key-chain bfd-pim key 53 secret
  $9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit protocols pim]
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication
  loose-check
```

5. (Optional) View your configuration by using the **show bfd session detail** or **show bfd session extensive** command.
6. Repeat these steps to configure the other end of the BFD session.

Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration by using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **ge-0/1/5** interface. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-pim**. The authentication keychain is configured with two keys. Key **1** contains the secret data "**9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm**" and a start time of June 1, 2009, at 9:46:02 AM PST. Key **2** contains the secret data "**9a5jiKW9l.reP38ny.TszF2/9**" and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols pim]
interface ge-0/1/5 {
  family inet {
    bfd-liveness-detection {
      authentication {
        key-chain bfd-pim;
        algorithm keyed-sha-1;
      }
    }
  }
}
[edit security]
authentication key-chains {
  key-chain bfd-pim {
    key 1 {
      secret "$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm";
      start-time "2009-6-1.09:46:02 -0700";
    }
```

```
        key 2 {
          secret "$9$a5jiKW9l.reP38ny.TszF2/9";
          start-time "2009-6-1.15:29:20 -0700";
        }
      }
    }
```

If you commit these updates to your configuration, you see output similar to the following example. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

show bfd session detail

```
user@host# show bfd session detail

                                                  Detect   Transmit
Address                 State     Interface       Time     Interval  Multiplier
50.0.0.2                Up        ge-0/1/5.0      0.900     0.300        3
 Client PIM, TX interval 0.300, RX interval 0.300, Authenticate
 Session up time 3d 00:34
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated
```

show bfd session extensive

```
user@host# show bfd session extensive
                                                  Detect   Transmit
Address                 State     Interface       Time     Interval  Multiplier
50.0.0.2                Up        ge-0/1/5.0      0.900     0.300        3
 Client PIM, TX interval 0.300, RX interval 0.300, Authenticate
        keychain bfd-pim, algo keyed-sha-1, mode strict
 Session up time 00:04:42
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated
 Min async interval 0.300, min slow interval 1.000
 Adaptive async TX interval 0.300, RX interval 0.300
 Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3
 Remote min TX interval 0.300, min RX interval 0.300, multiplier 3
 Local discriminator 2, remote discriminator 2
 Echo mode disabled/inactive
 Authentication enabled/active, keychain bfd-pim, algo keyed-sha-1, mode strict
```

Related Documentation

- *Understanding Bidirectional Forwarding Detection Authentication for PIM*
- Configuring BFD for PIM on page 107
- *authentication-key-chains*
- bfd-liveness-detection on page 182
- *show bfd session* in the CLI Explorer

CHAPTER 23

IGMP

Configuring IGMP

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.

6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements. See *Configuring the Session Announcement Protocol*.

To configure the Internet Group Management Protocol (IGMP), include the **igmp** statement:

```
igmp {
  accounting;
  interface interface-name {
    disable;
    (accounting | no-accounting);
    group-policy [ policy-names ];
    immediate-leave;
    oif-map map-name;
    promiscuous-mode;
    ssm-map ssm-map-name;
    static {
      group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
          source-count number;
          source-increment increment;
        }
      }
    }
    version version;
  }
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

By default, IGMP is enabled on all interfaces on which you configure Protocol Independent Multicast (PIM), and on all broadcast interfaces on which you configure the Distance Vector Multicast Routing Protocol (DVMRP).

NOTE: You can configure IGMP on an interface without configuring PIM. PIM is generally not needed on IGMP downstream interfaces. Therefore, only one "pseudo PIM interface" is created to represent all IGMP downstream (IGMP-only) interfaces on the router. This reduces the amount of router resources, such as memory, that are consumed. You must configure PIM on upstream IGMP interfaces to enable multicast routing, perform reverse-path forwarding for multicast data packets, populate the multicast forwarding table for upstream interfaces, and in the case of bidirectional PIM and PIM sparse mode, to distribute IGMP group memberships into the multicast routing domain.

Enabling IGMP

The Internet Group Management Protocol (IGMP) manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use IGMP to learn which groups have members on each of their attached physical networks. IGMP must be enabled for the router to receive IPv4 multicast packets. IGMP is only needed for IPv4 networks, because multicast is handled differently in IPv6 networks. IGMP is automatically enabled on all IPv4 interfaces on which you configure PIM and on all IPv4 broadcast interfaces when you configure DVMRP.

If IGMP is not running on an interface—either because PIM and DVMRP are not configured on the interface or because IGMP is explicitly disabled on the interface—you can explicitly enable IGMP.

To explicitly enable IGMP:

1. If PIM and DVMRP are not running on the interface, explicitly enable IGMP by including the interface name.

   ```
   [edit protocols igmp]
   user@host# set interface fe-0/0/0.0
   ```

2. See if IGMP is disabled on any interfaces. In the following example, IGMP is disabled on a Gigabit Ethernet interface.

   ```
   [edit protocols igmp]
   user@host# show
   interface fe-0/0/0.0;
   interface ge-1/0/0.0 {
     disable;
   }
   ```

3. Enable IGMP on the interface by deleting the **disable** statement.

   ```
   [edit protocols igmp]
   delete interface ge-1/0/0.0 disable
   ```

4. Verify the configuration.

   ```
   [edit protocols igmp]
   user@host# show
   interface fe-0/0/0.0;
   ```

```
interface ge-1/0/0.0;
```

5. Verify the operation of IGMP on the interfaces by checking the output of the **show igmp interface** command.

Related Documentation

- Understanding IGMP on page 32
- Disabling IGMP on page 134
- show igmp interface on page 357

Changing the IGMP Version

By default, the routing device runs IGMPv2. Routing devices running different versions of IGMP determine the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version.

To enable source-specific multicast (SSM) functionality, you must configure version 3 on the host and the host's directly connected routing device. If a source address is specified in a multicast group that is statically configured, the version must be set to IGMPv3.

If a static multicast group is configured with the source address defined, and the IGMP version is configured to be version 2, the source is ignored and only the group is added. In this case, the join is treated as an IGMPv2 group join.

If you configure the IGMP version setting at the individual interface hierarchy level, it overrides the **interface all** statement.

If you have already configured the routing device to use IGMP version 1 (IGMPv1) and then configure it to use IGMPv2, the routing device continues to use IGMPv1 for up to 6 minutes and then uses IGMPv2.

To change to IGMPv3 for SSM functionality:

1. Configure the IGMP interface.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0 version 3
```

2. Verify the configuration by checking the version field in the output of the **show igmp interfaces** command. The **show igmp statistics** command has version-specific output fields, such as V1 Membership Report, V2 Membership Report, and V3 Membership Report.

Related Documentation

- Understanding IGMP on page 32
- show pim interfaces on page 417
- show igmp statistics on page 361
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 3376, *Internet Group Management Protocol, Version 3*

Modifying the IGMP Host-Query Message Interval

The objective of IGMP is to keep routers up to date with group membership of the entire subnet. Routers need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one router is elected the querier. The IGMP querier router periodically sends general host-query messages on each attached network to solicit membership information. The messages are sent to the all-systems multicast group address, 224.0.0.1.

The query interval, the response interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

By default, host-query messages are sent every 125 seconds. You can change this interval to change the number of IGMP messages sent on the subnet.

To modify the query interval:

1. Configure the interval.

   ```
   [edit protocols igmp]
   user@host# set query-interval 200
   ```

   The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the IGMP Query Interval field in the output of the **show igmp interface** command.

3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

Related Documentation

- Understanding IGMP on page 32
- Modifying the IGMP Query Response Interval on page 121
- Modifying the IGMP Robustness Variable on page 122
- **show igmp interface on page 357**
- **show igmp statistics on page 361**

Modifying the IGMP Last-Member Query Interval

The last-member query interval is the maximum amount of time between group-specific query messages, including those sent in response to leave-group messages. You can configure this interval to change the amount of time it takes a routing device to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group message from a host, the routing device sends multiple group-specific queries to the group being left. The querier sends a specific number of these queries at a specific interval. The number of queries sent is called the last-member query count. The interval at which the queries are sent is called the last-member query interval. Because both settings are configurable, you can adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

The last-member query count x (times) the last-member query interval = (equals) the amount of time it takes a routing device to determine that the last member of a group has left the group and to stop forwarding group traffic.

The default last-member query interval is 1 second. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify this interval:

1. Configure the time (in seconds) that the routing device waits for a report in response to a group-specific query.

   ```
   [edit protocols igmp]
   user@host# set query-last-member-interval 0.1
   ```

2. Verify the configuration by checking the IGMP Last Member Query Interval field in the output of the **show igmp interfaces** command.

NOTE: You can configure the last-member query count by configuring the robustness variable. The two are always equal.

Related Documentation

- Modifying the IGMP Robustness Variable on page 122
- show pim interfaces on page 417

Specifying Immediate-Leave Host Removal for IGMP

The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.

NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

To enable immediate leave on an interface:

1. Configure immediate leave on the IGMP interface.

   ```
   [edit protocols IGMP]
   user@host# set interface ge-0/0/0.1 immediate-leave
   ```

2. Verify the configuration by checking the Immediate Leave field in the output of the **show igmp interface** command.

Related Documentation

- Understanding IGMP on page 32
- show igmp interface on page 357

Filtering Unwanted IGMP Reports at the IGMP Interface Level

Suppose you need to limit the subnets that can join a certain multicast group. The **group-policy** statement enables you to filter unwanted IGMP reports at the interface level. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only IGMP group addresses (for IGMPv2) by using the policy's **route-filter** statement to match the group address. You define the policy to match IGMP (source, group) addresses (for IGMPv3) by using the policy's **route-filter** statement to match the group address and the policy's **source-address-filter** statement to match the source address.

To filter unwanted IGMP reports:

1. Configure an IGMPv2 policy.

   ```
   [edit policy-statement reject_policy_v2]
   user@host# set from route-filter 224.1.1.1/32 exact
   ```

```
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set then reject
```

2. Configure an IGMPv3 policy.

```
[edit policy-statement reject_policy_v3]
user@host# set from route-filter 224.1.1.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject
```

3. Apply the policies to the IGMP interfaces on which you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running IGMPv2, and **ge-0/1/1.0** is running IGMPv3.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v2
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v3
```

4. Verify the operation of the filter by checking the Rejected Report field in the output of the **show igmp statistics** command.

Related Documentation

- Understanding IGMP on page 32
- *Example: Configuring Policy Chains and Route Filters*
- show igmp statistics on page 361

Accepting IGMP Messages from Remote Subnetworks

By default, IGMP interfaces accept IGMP messages only from the same subnet. Including the **promiscuous-mode** statement enables the routing device to accept IGMP messages from indirectly connected subnets.

NOTE: When you enable IGMP on an unnumbered Ethernet interface that uses a /32 loopback address as a donor address, you must configure IGMP promiscuous mode to accept the IGMP packets received on this interface.

NOTE: When enabling promiscuous-mode, all routers on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.

To enable IGMP promiscuous mode on an interface:

1. Configure the IGMP interface.

```
[edit protocols igmp]
user@host# set interface ge-0/1/1.0 promiscuous-mode
```

2. Verify the configuration by checking the Promiscuous Mode field in the output of the **show igmp interface** command.
3. Verify the operation of the filter by checking the Rx non-local field in the output of the **show igmp statistics** command.

Related Documentation

- Understanding IGMP on page 32
- *Configuring the Loopback Interface* in the *Junos OS Network Interfaces Library for Routing Devices*
- show igmp interface on page 357
- show igmp statistics on page 361

Modifying the IGMP Query Response Interval

The query response interval is the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. Configuring this interval allows you to adjust the burst peaks of IGMP messages on the subnet. Set a larger interval to make the traffic less bursty. Bursty traffic refers to an uneven pattern of data transmission: sometimes a very high data transmission rate, whereas at other times a very low data transmission rate.

The query response interval, the host-query interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely originated multicast packets for that group onto the attached network.

The default query response interval is 10 seconds. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify the query response interval:

1. Configure the interval.

```
[edit protocols igmp]
user@host# set query-response-interval 0.4
```

2. Verify the configuration by checking the IGMP Query Response Interval field in the output of the **show igmp interface** command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

Related Documentation

- Understanding IGMP on page 32
- Modifying the IGMP Host-Query Message Interval on page 117

- Modifying the IGMP Robustness Variable on page 122
- show igmp interface on page 357
- show igmp statistics on page 361

Modifying the IGMP Robustness Variable

Fine-tune the IGMP robustness variable to allow for expected packet loss on a subnet. The robust count automatically changes certain IGMP message intervals for IGMPv2 and IGMPv3. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

When the query router receives an IGMP leave message on a shared network running IGMPv2, the query router must send an IGMP group query message a specified number of times. The number of IGMP group query messages sent is determined by the robust count.

The value of the robustness variable is also used in calculating the following IGMP message intervals:

- Group member interval—Amount of time that must pass before a multicast router determines that there are no more members of a group on a network. This interval is calculated as follows: (robustness variable x query-interval) + (1 x query-response-interval).
- Other querier present interval—The robust count is used to calculate the amount of time that must pass before a multicast router determines that there is no longer another multicast router that is the querier. This interval is calculated as follows: (robustness variable x query-interval) + (0.5 x query-response-interval).
- Last-member query count—Number of group-specific queries sent before the router assumes there are no local members of a group. The number of queries is equal to the value of the robustness variable.

In IGMPv3, a change of interface state causes the system to immediately transmit a state-change report from that interface. In case the state-change report is missed by one or more multicast routers, it is retransmitted. The number of times it is retransmitted is the robust count minus one. In IGMPv3, the robust count is also a factor in determining the group membership interval, the older version querier interval, and the other querier present interval.

By default, the robustness variable is set to 2. You might want to increase this value if you expect a subnet to lose packets.

The number can be from 2 through 10.

To change the value of the robustness variable:

1. Configure the robust count.

   When you set the robust count, you are in effect configuring the number of times the querier retries queries on the connected subnets.

```
[edit protocols igmp]
user@host# set robust-count 5
```

2. Verify the configuration by checking the IGMP Robustness Count field in the output of the **show igmp interfaces** command.

Related Documentation

- Modifying the IGMP Host-Query Message Interval on page 117
- Modifying the IGMP Query Response Interval on page 121
- Modifying the IGMP Last-Member Query Interval on page 117
- show pim interfaces on page 417
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 3376, *Internet Group Management Protocol, Version 3*

Limiting the Maximum IGMP Message Rate

This section describes how to change the limit for the maximum number of IGMP packets transmitted in 1 second by the router.

Increasing the maximum number of IGMP packets transmitted per second might be useful on a router with a large number of interfaces participating in IGMP.

To change the limit for the maximum number of IGMP packets the router can transmit in 1 second, include the **maximum-transmit-rate** statement and specify the maximum number of packets per second to be transmitted.

Related Documentation

- maximum-transmit-rate (Protocols IGMP) on page 257

Enabling IGMP Static Group Membership

You can create IGMP static group membership to test multicast forwarding without a receiver host. When you enable IGMP static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts. The router on which you enable static IGMP group membership must be the designated router (DR) for the subnet. Otherwise, traffic does not flow downstream.

When enabling IGMP static group membership, you cannot configure multiple groups using the **group-count**, **group-increment**, **source-count**, and **source-increment** statements if the **all** option is specified as the IGMP interface.

Class-of-service (CoS) adjustment is not supported with IGMP static group membership.

In this example, you create static group 225.1.1.1.

1. On the DR, configure the static groups to be created by including the **static** statement and **group** statement and specifying which IP multicast address of the group to be created. When creating groups individually, you must specify a unique address for each group.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  static {
    group 225.1.1.1 ;
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created.

```
user@host> show igmp group
Interface: fe-0/1/2
     Group: 225.1.1.1
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
```

NOTE: When you configure static IGMP group entries on point-to-point links that connect routing devices to a rendezvous point (RP), the static IGMP group entries do not generate join messages toward the RP.

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of static groups be automatically created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately.

In this example, you create three groups.

1. On the DR, configure the number of static groups to be created by including the **group-count** statement and specifying the number of groups to be created.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  static {
    group 225.1.1.1 {
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.2, and 225.1.1.3 have been created.

```
user@host> show igmp group
Interface: fe-0/1/2
     Group: 225.1.1.1
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
     Group: 225.1.1.2
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
     Group: 225.1.1.3
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can also configure the group address to be automatically incremented for each group created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and when you do not want the group addresses to be sequential.

In this example, you create three groups and increase the group address by an increment of two for each group.

1. On the DR, configure the group address increment by including the **group-increment** statement and specifying the number by which the address should be incremented for each group. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3 group-increment
  0.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      group-increment 0.0.0.2;
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.3, and 225.1.1.5 have been created.

```
user@host> show igmp group
```

```
Interface: fe-0/1/2
     Group: 225.1.1.1
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
     Group: 225.1.1.3
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
     Group: 225.1.1.5
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, and your network is operating in source-specific multicast (SSM) mode, you can also specify that the multicast source address be accepted. This is useful when you want to test forwarding to multicast receivers from a specific multicast source.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you create group 225.1.1.1 and accept IP address 10.0.0.2 as the only source.

1. On the DR, configure the source address by including the **source** statement and specifying the IPv4 address of the source host.

   ```
   [edit protocols igmp]
   user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2
   ```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

   ```
   user@host> show configuration protocol igmp
   ```

   ```
   interface fe-0/1/2.0 {
     version 3;
     static {
       group 225.1.1.1 {
         source 10.0.0.2;
       }
     }
   }
   ```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that source 10.0.0.2 has been accepted.

   ```
   user@host> show igmp group
   Interface: fe-0/1/2
        Group: 225.1.1.1
            Source: 10.0.0.2
            Last reported by: Local
            Timeout: 0 Type: Static
   ```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of multicast sources be automatically accepted. This is useful when you want to test forwarding to multicast receivers from more than one specified multicast source.

In this example, you create group 255.1.1.1 and accept addresses 10.0.0.2, 10.0.0.3, and 10.0.0.4 as the sources.

1. On the DR, configure the number of multicast source addresses to be accepted by including the **source-count** statement and specifying the number of sources to be accepted.

   ```
   [edit protocols igmp]
   user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count
     3
   ```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

   ```
   user@host> show configuration protocol igmp

   interface fe-0/1/2.0 {
     version 3;
     static {
       group 225.1.1.1 {
         source 10.0.0.2 {
           source-count 3;
         }
       }
     }
   }
   ```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.3, and 10.0.0.4 have been accepted.

   ```
   user@host> show igmp group
   Interface: fe-0/1/2
        Group: 225.1.1.1
            Source: 10.0.0.2
            Last reported by: Local
            Timeout: 0 Type: Static
        Group: 225.1.1.1
            Source: 10.0.0.3
            Last reported by: Local
            Timeout: 0 Type: Static
        Group: 225.1.1.1
            Source: 10.0.0.4
            Last reported by: Local
            Timeout: 0 Type: Static
   ```

When you configure static groups on an interface on which you want to receive multicast traffic, and specify that a number of multicast sources be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and you do not want the source addresses to be sequential.

In this example, you create group 225.1.1.1 and accept addresses 10.0.0.2, 10.0.0.4, and 10.0.0.6 as the sources.

1. Configure the multicast source address increment by including the **source-increment** statement and specifying the number by which the address should be incremented for each source. The increment is specified in dotted decimal notation similar to an IPv4 address.

   ```
   [edit protocols igmp]
   user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count
     3 source-increment 0.0.0.2
   ```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

   ```
   user@host> show configuration protocol igmp

   interface fe-0/1/2.0 {
     version 3;
     static {
       group 225.1.1.1 {
         source 10.0.0.2 {
           source-count 3;
           source-increment 0.0.0.2;
         }
       }
     }
   }
   ```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.4, and 10.0.0.6 have been accepted.

   ```
   user@host> show igmp group
   Interface: fe-0/1/2
        Group: 225.1.1.1
            Source: 10.0.0.2
            Last reported by: Local
            Timeout: 0 Type: Static
        Group: 225.1.1.1
            Source: 10.0.0.4
            Last reported by: Local
            Timeout: 0 Type: Static
        Group: 225.1.1.1
            Source: 10.0.0.6
            Last reported by: Local
            Timeout: 0 Type: Static
   ```

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the source address configured. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the source address configured.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you exclude address 10.0.0.2 as a source for group 225.1.1.1.

1. On the DR, configure a multicast static group to operate in exclude mode by including the **exclude** statement and specifying which IPv4 source address to exclude.

   ```
   [edit protocols igmp]
   user@host# set interface fe-0/1/2 static group 225.1.1.1 exclude source 10.0.0.2
   ```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

   ```
   user@host> show configuration protocol igmp

   interface fe-0/1/2.0 {
     version 3;
     static {
       group 225.1.1.1 {
         exclude;
         source 10.0.0.2;
       }
     }
   }
   ```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group detail** command to verify that static group 225.1.1.1 has been created and that the static group is operating in exclude mode.

   ```
   user@host> show igmp group detail
   Interface: fe-0/1/2
        Group: 225.1.1.1
            Group mode: Exclude
            Source: 10.0.0.2
            Last reported by: Local
            Timeout: 0 Type: Static
   ```

Related Documentation

- *Enabling MLD Static Group Membership*
- group (Protocols IGMP) on page 249
- group-count (Protocols IGMP) on page 250
- group-increment (Protocols IGMP) on page 250
- source-count (Protocols IGMP) on page 264

- source-increment (Protocols IGMP) on page 264
- static (Protocols IGMP) on page 266

Recording IGMP Join and Leave Events

To determine whether IGMP tuning is needed in a network, you can configure the routing device to record IGMP join and leave events. You can record events globally for the routing device or for individual interfaces.

Table 5 on page 130 describes the recordable IGMP events.

Table 5: IGMP Event Messages

ERRMSG Tag	Definition
RPD_IGMP_JOIN	Records IGMP join events.
RPD_IGMP_LEAVE	Records IGMP leave events.
RPD_IGMP_ACCOUNTING_ON	Records when IGMP accounting is enabled on an IGMP interface.
RPD_IGMP_ACCOUNTING_OFF	Records when IGMP accounting is disabled on an IGMP interface.
RPD_IGMP_MEMBERSHIP_TIMEOUT	Records IGMP membership timeout events.

To enable IGMP accounting:

1. Enable accounting globally or on an IGMP interface. This example shows both options.

   ```
   [edit protocols igmp]
   user@host# set accounting
   user@host# set interface fe-0/1/0.2 accounting
   ```

2. Configure the events to be recorded and filter the events to a system log file with a descriptive filename, such as **igmp-events**.

   ```
   [edit system syslog file igmp-events]
   user@host# set any info
   user@host# set match ".*RPD_IGMP_JOIN.* | .*RPD_IGMP_LEAVE.* |
     .*RPD_IGMP_ACCOUNTING.* | .*RPD_IGMP_MEMBERSHIP_TIMEOUT.*"
   ```

3. Periodically archive the log file.

   This example rotates the file size when it reaches 100 KB and keeps three files.

   ```
   [edit system syslog file igmp-events]
   user@host# set archive size 100000
   user@host# set archive files 3
   user@host# set archive archive-sites "ftp://user@host1//var/tmp" password
     "anonymous"
   user@host# set archive archive-sites "ftp://user@host2//var/tmp" password "test"
   user@host# set archive transfer-interval 24
   user@host# set archive start-time 2011-01-07:12:30
   ```

4. You can monitor the system log file as entries are added to the file by running the ***monitor start*** and **monitor stop** commands.

   user@host> **monitor start igmp-events**

```
*** igmp-events ***
Apr 16 13:08:23  host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command
'run monitor start igmp-events '
monitor
```

Related Documentation

- Understanding IGMP on page 32
- *Specifying Log File Size, Number, and Archiving Properties*

Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

The **group-limit** statement enables you to limit the number of IGMP multicast group joins for logical interfaces. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for IGMP multicast groups, keep the following in mind:

- Each any-source group (*,G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in IGMPv3 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The groups must then request to rejoin the network (up to the newly configured group limit).
- You can dynamically limit multicast groups on IGMP logical interfaces using dynamic profiles.

Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for IGMP multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of IGMP multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The **group-threshold** statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a

percentage of the group limit, so you must configure the **group-limit** statement to configure a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs the warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the **log-interval** statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for IGMP multicast group joins.

To limit multicast group joins on an IGMP logical interface:

1. Access the logical interface at the IGMP protocol hierarchy level.

   [edit]
   user@host# **edit protocols igmp interface** ***interface-name***

2. Specify the group limit for the interface.

   [edit protocols igmp interface *interface-name*]
   user@host# **set** group-limit ***limit***

3. (Optional) Configure the threshold at which a warning message is logged.

   [edit protocols igmp interface *interface-name*]
   user@host# **set group-threshold** ***value***

4. (Optional) Configure the amount of time between log messages.

   [edit protocols igmp interface *interface-name*]
   user@host# **set log-interval** ***seconds***

To confirm your configuration, use the **show protocols igmp** command. To verify the operation of IGMP on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the show igmp interface command.

Related Documentation

- Enabling IGMP Static Group Membership on page 123

Tracing IGMP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
client-notification	Trace notifications.

Flag	Description
general	Trace general flow.
group	Trace group operations.
host-notification	Trace host notifications.
leave	Trace leave group messages (IGMPv2 only).
mtrace	Trace mtrace packets. Use the **mtrace** command to troubleshoot the software.
normal	Trace normal events.
packets	Trace all IGMP packets.
policy	Trace policy processing.
query	Trace IGMP membership query messages, including general and group-specific queries.
report	Trace membership report messages.
route	Trace routing information.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on IGMP packets of a particular type. To configure tracing operations for IGMP:

1. (Optional) Configure tracing at the routing options level to trace all protocol packets.

   ```
   [edit routing-options traceoptions]
   user@host# set file all-packets-trace
   user@host# set flag all
   ```

2. Configure the filename for the IGMP trace file.

   ```
   [edit protocols igmp traceoptions]
   user@host# set file igmp-trace
   ```

3. (Optional) Configure the maximum number of trace files.

   ```
   [edit protocols igmp traceoptions]
   user@host# set file files 5
   ```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols igmp traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols igmp traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with a particular multicast group. The following example shows how to flag all events for packets associated with the group IP address.

```
[edit protocols igmp traceoptions]
user@host# set flag group | match 232.1.1.2
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/igmp-trace
```

Related Documentation

- Understanding IGMP on page 32
- *Junos OS Tracing and Logging Operations* in the *Junos OS Administration Library for Routing Devices*
- **mtrace on page 340** in the CLI Explorer

Disabling IGMP

To disable IGMP on an interface, include the **disable** statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols igmp interface *interface-name*]**
- **[edit logical-systems *logical-system-name* protocols igmp interface *interface-name*]**

Related Documentation

- Enabling IGMP on page 115

CHAPTER 24

IGMP Snooping

Configuring IGMP Snooping

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).

NOTE: You cannot configure IGMP snooping on a secondary VLAN.

To enable IGMP snooping and configure individual options as needed for your network by using the CLI:

1. Enable IGMP snooping on a VLAN:

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan employee-vlan
   ```

2. Configure the switch to immediately remove group membership from interfaces on a VLAN when it receives a leave message through that VLAN, and have it not forward any membership queries for the multicast group to the VLAN (IGMPv2 only):

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan vlan-name immediate-leave
   ```

3. Configure an interface to belong to a multicast group:

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan-name interface interface-name static group
   group-address
   ```

4. Configure an interface to forward IGMP queries received from multicast routers.

   ```
   [edit protocols]
   ```

```
user@switch# set igmp-snooping vlan vlan-name interface interface-name
multicast-router-interface
```

5. Configure the switch to wait for four timeout intervals before timing out a multicast group on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name robust-count 4
```

6. If you want the switch to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name igmp-querier source-address source address
```

The switch uses the address that you configure as the source address in the IGMP queries that it sends. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is greater (a higher number) than the IP addresses for those routers on the network. This ensures that switch is always the IGMP querier on the network.

NOTE: The igmp-querier statement is not supported on QFabric systems.

Related Documentation

- IGMP Snooping Overview on page 35
- Example: Configuring IGMP Snooping on page 137
- Changing the IGMP Snooping Group Timeout Value on page 136
- Monitoring IGMP Snooping on page 315

Changing the IGMP Snooping Group Timeout Value

The IGMP snooping group timeout value determines how long a switch waits to receive an IGMP query from a multicast router before removing a multicast group from its multicast cache table. A switch calculates the timeout value by using the **query-interval** and **query-response-interval** values.

When you enable IGMP snooping, the **query-interval** and **query-response-interval** values are applied to all VLANs on the switch. The values are:

- **query-interval**—125 seconds
- **query-response-interval**—10 seconds

The switch automatically calculates the group timeout value for an IGMP snooping-enabled switch by multiplying the **query-interval** value by 2 (the default **robust-count** value) and then adding the **query-response-interval** value. By default, the switch waits 260 seconds to receive an IGMP query before removing a multicast group from its multicast cache table: (125 x 2) + 10 = 260.

You can modify the group timeout value by changing the **robust-count** value. For example, if you want the system to wait 510 seconds before timing groups out—(125 x 4) + 10 = 510—enter this command:

```
[edit protocols]
```

```
user@switch# set igmp-snooping vlan employee-vlan robust-count (IGMP Snooping) 4
```

Related Documentation

- Verifying the IGMP Snooping Group Timeout Value on page 316
- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135

Example: Configuring IGMP Snooping

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).

This example describes how to configure IGMP snooping:

Requirements

This example requires Junos OS Release 11.1 or later on a QFX Series product.

Before you configure IGMP snooping, be sure you have:

- Configured the **employee-vlan** VLAN
- Assigned interfaces **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/3** to **employee-vlan**

Overview and Topology

In this example you configure an interface to receive multicast traffic from a source and configure some multicast-related behavior for downstream interfaces. The example assumes that IGMP snooping was previously disabled for the VLAN.

Table 6 on page 137 shows the components of the topology for this example.

Table 6: Components of the IGMP Snooping Topology

Components	Settings
VLAN name	**employee-vlan**, tag **20**
Interfaces in **employee-vlan**	**ge-0/0/1**, **ge-0/0/2**, **ge-0/0/3**
Multicast IP address for **employee-vlan**	**225.100.100.100**

Configuration

To configure basic IGMP snooping on a switch:

CLI Quick Configuration

To quickly configure IGMP snooping, copy the following commands and paste them into a terminal window:

```
[edit protocols]
set igmp-snooping vlan employee-vlan
set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group 225.100.100.100
set igmp-snooping vlan employee-vlan interface ge-0/0/2 multicast-router-interface
set igmp-snooping vlan employee-vlan robust-count 4
```

Step-by-Step Procedure

Configure IGMP snooping:

1. Enable and configure IGMP snooping on the VLAN **employee-vlan**:

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan employee-vlan
   ```

2. Configure a interface to belong to a multicast group:

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group
   225.100.100.100
   ```

3. Configure an interface to forward IGMP queries received from multicast routers.

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/2
   multicast-router-interface
   ```

4. Configure the switch to wait for four timeout intervals before timing out a multicast group on a VLAN:

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan employee-vlan robust-count 4
   ```

Results

Check the results of the configuration:

```
user@switch# show protocols igmp-snooping
  vlan employee-vlan {
    robust-count 4;
    }
    interface ge-0/0/2 {
      multicast-router-interface;
    }
    interface ge-0/0/3 {
      static {
        group 255.100.100.100;
      }
    }
  }
```

Related Documentation

- IGMP Snooping Overview on page 35
- Configuring IGMP Snooping on page 135
- Changing the IGMP Snooping Group Timeout Value on page 136
- Monitoring IGMP Snooping on page 315

- *Example: Setting Up Bridging with Multiple VLANs*.

Configuring Multicast VLAN Registration (CLI Procedure)

Multicast VLAN registration (MVR) enables hosts that are not part of a multicast source VLAN (MVLAN) to still receive multicast streams from the MVLAN, allowing an MVLAN to be shared across a Layer 2 network. Hosts remain in their own VLANs for bandwidth and security reasons but are able to receive multicast streams from the MVLAN.

You can configure one or more VLANs on a switch to be MVLANs or MVR receiver VLANs. By default, MVR is not configured on EX Series switches and the QFX Series.

NOTE: MVR is supported on VLANs running IGMP version 2 (IGMPv2) only.

NOTE: When you configure MVR, the following restrictions apply:

- You cannot enable multicast protocols on VLAN interfaces that are members of MVLANs.
- If you configure an MVLAN in proxy mode, IGMP snooping proxy mode is automatically enabled on all MVR receiver VLANs of this MVLAN. If a VLAN is an MVR receiver VLAN for multiple MVLANs, all of the MVLANs must have proxy mode enabled or all must have proxy mode disabled. You can enable proxy mode only on VLANs that are configured as MVR source VLANs and that are not configured for Q-in-Q tunneling.
- After you configure a VLAN as an MVLAN, that VLAN is no longer available for other uses.

To configure MVR:

1. Configure the VLAN named mv0 to be an MVLAN:

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan mv0 data-forwarding source groups
     225.10.0.0/16
   ```

2. Configure the MVLAN mv0 to be a proxy VLAN:

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan mv0 proxy source-address 10.0.0.1
   ```

3. Configure the VLAN named v2 to be an MVR receiver VLAN with mv0 as its source:

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan v2 data-forwarding receiver source-vlans mv0
   ```

4. Install forwarding entries in the MVR receiver VLAN:

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan mv0 data-forwarding receiver install
   ```

Related Documentation

- Example: Configuring Multicast VLAN Registration on page 140
- Understanding Multicast VLAN Registration on page 47

Example: Configuring Multicast VLAN Registration

Multicast VLAN registration (MVR) enables hosts that are not part of a multicast VLAN (MVLAN) to receive multicast streams from the MVLAN, which enable the MVLAN to be shared across the Layer 2 network and eliminate the need to send duplicate multicast streams to each requesting VLAN in the network. Hosts remain in their own VLANs for bandwidth and security reasons.

This example describes how to configure MVR on EX Series switches and the QFX Series.

Requirements

This example uses the following hardware and software components:

- One EX Series switch or the QFX Series
- Junos OS Release 9.6 or later for EX Series switches or Junos OS Release 12.3 or later for the QFX Series

Before you configure MVR, be sure you have:

- Configured two or more VLANs on the switch. See the task for your platform:
  - *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
  - *Example: Setting Up Bridging with Multiple VLANs* for the QFX Series
- Connected the switch to a network that can transmit IPTV multicast streams from a video server.
- Connected a host that is capable of receiving IPTV multicast streams to an interface in one of the VLANs.

Overview and Topology

In a standard Layer 2 network, a multicast stream received on one VLAN is never distributed to interfaces outside that VLAN. If hosts in multiple VLANs request the same multicast stream, a separate copy of that multicast stream is distributed to the requesting VLANs.

MVR introduces the concept of a *multicast source VLAN* (MVLAN), which is created by MVR and becomes the only VLAN over which multicast traffic flows throughout the Layer 2 network. Multicast traffic can then be selectively forwarded from interfaces on the MVLAN (source ports) to hosts that are connected to interfaces (multicast receiver

ports) that are not part of the multicast source VLAN. When you configure an MVLAN, you assign a range of multicast group addresses to it. You then configure other VLANs to be MVR receiver VLANs, which receive multicast streams from the MVLAN. The MVR receiver ports comprise all the interfaces that exist on any of the MVR receiver VLANs.

You can configure MVR to operate in one of two modes: transparent mode (the default mode) or proxy mode. Both modes enable MVR to forward only one copy of a multicast stream to the Layer 2 network.

In transparent mode, the switch receives one copy of each IPTV multicast stream and then replicates the stream only to those hosts that want to receive it, while forwarding all other types of multicast traffic without modification. Figure 16 on page 142 shows how MVR operates in transparent mode.

In proxy mode, the switch acts as a proxy for the IGMP multicast router in the MVLAN for MVR group memberships established in the MVR receiver VLANs and generates and sends IGMP packets into the MVLAN as needed. Figure 17 on page 143 shows how MVR operates in proxy mode.

This example shows how to configure MVR in both transparent mode and proxy mode on an EX Series switch or the QFX Series. The topology includes a video server that is connected to a multicast router, which in turn forwards the IPTV multicast traffic in the MVLAN to the Layer 2 network.

Figure 16 on page 142 shows the MVR topology in transparent mode. Interfaces P1 and P2 on Switch C belong to service VLAN s0 and MVLAN mv0. Interface P4 of Switch C also belongs to service VLAN s0. In the upstream direction of the network, only non-IPTV traffic is being carried in individual customer VLANs of service VLAN s0. VLAN c0 is an example of this type of customer VLAN. IPTV traffic is being carried on MVLAN mv0. If any host on any customer VLAN connected to port P4 requests an MVR stream, Switch C takes the stream from VLAN mv0 and replicates that stream onto port P4 with tag mv0. IPTV traffic, along with other network traffic, flows from port P4 out to the Digital Subscriber Line Access Multiplexer (DSLAM) D1.

Figure 16: MVR Topology in Transparent Mode

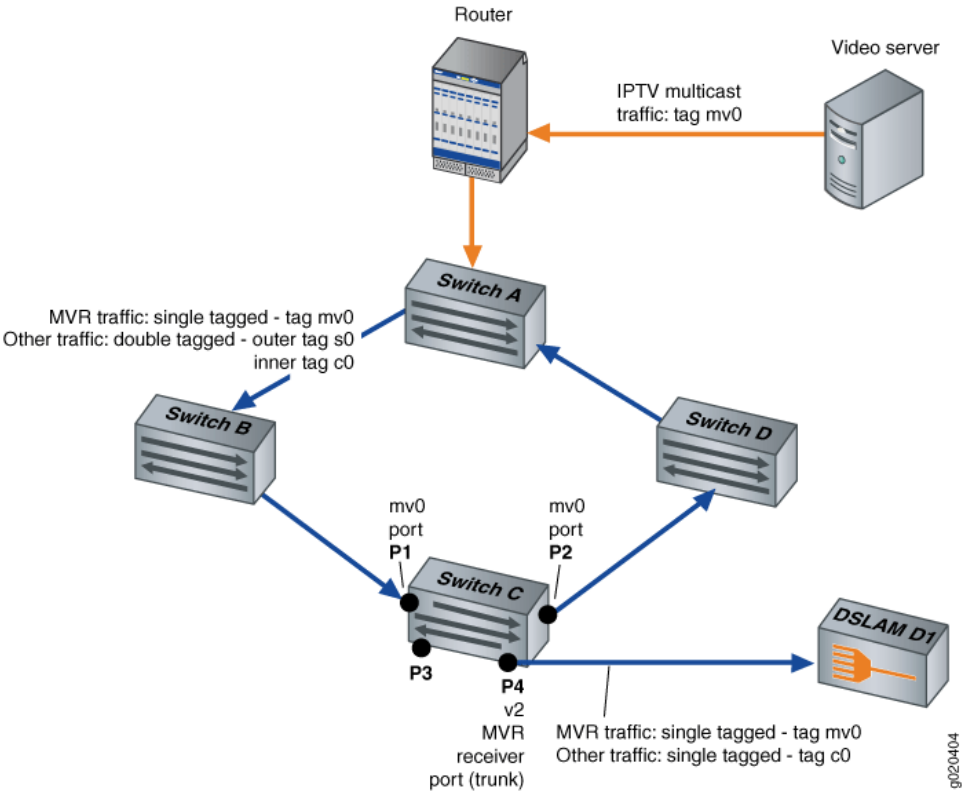

Figure 17 on page 143 shows the MVR topology in proxy mode. Interfaces P1 and P2 on Switch C belong to MVLAN mv0 and customer VLAN c0. Interface P4 on Switch C is an access port of customer VLAN c0. In the upstream direction of the network, only non-IPTV traffic is being carried on customer VLAN c0. Any IPTV traffic requested by hosts on VLAN c0 is replicated untagged to port P4 based on streams received in MVLAN mv0. IPTV traffic flows from port P4 out to an IPTV-enabled device in Host H1. Other traffic, such as data and voice traffic, also flows from port P4 to other network devices in Host H1.

Figure 17: MVR Topology in Proxy Mode

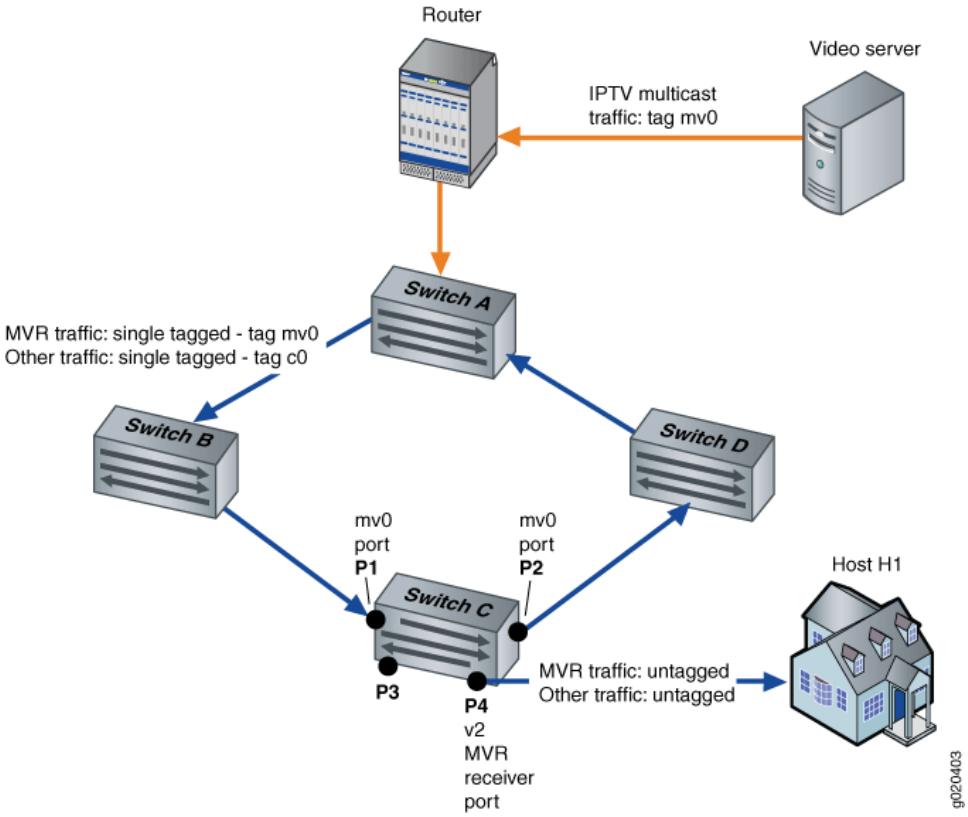

For information on VLAN tagging, see the topic for your platform:

- *Understanding Bridging and VLANs on EX Series Switches*
- *Understanding Bridging and VLANs* on the QFX Series

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit protocols igmp-snooping]** hierarchy level.

```
set vlan mv0 data-forwarding source groups 225.10.0.0/16
set vlan v2 data-forwarding receiver source-vlans mv0
set vlan v2 data-forwarding receiver install
set vlan mv0 proxy source-address 10.1.1.1
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure MVR:

1. Configure VLAN mv0 to be an MVLAN:

   ```
   [edit protocols igmp-snooping]
   user@switch# set vlan mv0 data-forwarding source groups 225.10.0.0/16
   ```

2. Configure VLAN v2 to be a multicast receiver VLAN with mv0 as its source:

   ```
   [edit protocols igmp-snooping]
   user@switch# set vlan v2 data-forwarding receiver source-vlans mv0
   ```

3. (Optional) Install forwarding entries in the multicast receiver VLAN v2:

   ```
   [edit protocols igmp-snooping]
   user@switch# set vlan v2 data-forwarding receiver install
   ```

4. (Optional) Configure MVR in proxy mode:

   ```
   [edit protocols igmp-snooping]
   user@switch# set vlan mv0 proxy source-address 10.1.1.1
   ```

Results

From configuration mode, confirm your configuration by entering the **show** command at the **[edit protocols igmp-snooping]** hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit protocols igmp-snooping]
user@switch# show
vlan mv0 {
  proxy {
    source-address 10.1.1.1;
  }
  data-forwarding {
    source {
      groups 225.10.0.0/16;
    }
  }
}
vlan v2 {
  data-forwarding {
    receiver {
      source-vlans mv0;
      install;
    }
  }
}
```

Related Documentation

- Configuring Multicast VLAN Registration (CLI Procedure) on page 139
- Understanding Multicast VLAN Registration on page 47

CHAPTER 25

MSDP

Configuring MSDP

To configure the Multicast Source Discovery Protocol (MSDP), include the **msdp** statement:

```
msdp {
  disable;
  active-source-limit {
    maximum number;
    threshold number;
  }
  data-encapsulation (disable | enable);
  export [ policy-names ];
  group group-name {
    ... group-configuration ...
  }
  hold-time seconds;
  import [ policy-names ];
  local-address address;
  keep-alive seconds;
  peer address {
    ... peer-configuration ...
  }
  rib-group group-name;
  source ip-prefix</prefix-length> {
    active-source-limit {
      maximum number;
      threshold number;
    }
  }
```

```
    sa-hold-time seconds;
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier > <disable>;
    }
    group group-name {
      disable;
      export [ policy-names ];
      import [ policy-names ];
      local-address address;
      mode (mesh-group | standard);
      peer address {
        ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
            just following ...
      }
      traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
      }
    }
    peer address {
      disable;
      active-source-limit {
        maximum number;
        threshold number;
      }
      authentication-key peer-key;
      default-peer;
      export [ policy-names ];
      import [ policy-names ];
      local-address address;
      traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
      }
    }
  }
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit routing-instances *routing-instance-name* protocols]**
- **[edit logical-systems *logical-system-name* protocols]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]**

By default, MSDP is disabled.

Related Documentation

- *Example: Configuring MSDP in a Routing Instance*
- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 150

Tracing MSDP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
general	Trace general events.
keepalive	Trace keepalive messages.
normal	Trace normal events.
packets	Trace all MSDP packets.
policy	Trace policy processing.
route	Trace MSDP changes to the routing table.
source-active	Trace source-active packets.
source-active-request	Trace source-active request packets.
source-active-response	Trace source-active response packets.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

You can configure MSDP tracing for all peers, for all peers in a particular group, or for a particular peer.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on MSDP peers in a particular group. To configure tracing operations for MSDP:

1. (Optional) Configure tracing by including the **traceoptions** statement at the **[edit routing-options]** hierarchy level and set the **all-packets-trace** and **all** flags to trace all protocol packets.

   ```
   [edit routing-options traceoptions]
   user@host# set file all-packets-trace
   user@host# set flag all
   ```

2. Configure the filename for the MSDP trace file.

   ```
   [edit protocols msdp group groupa traceoptions]
   user@host# set file msdp-trace
   ```

3. (Optional) Configure the maximum number of trace files.

   ```
   [edit protocols msdp group groupa traceoptions]
   user@host# set file files 5
   ```

4. (Optional) Configure the maximum size of each trace file.

   ```
   [edit protocols msdp group groupa traceoptions]
   user@host# set file size 1m
   ```

5. (Optional) Enable unrestricted file access.

   ```
   [edit protocols msdp group groupa traceoptions]
   user@host# set file world-readable
   ```

6. Configure tracing flags. Suppose you are troubleshooting issues with the source-active cache for **groupa**. The following example shows how to trace messages associated with the group address.

   ```
   [edit protocols msdp group groupa traceoptions]
   user@host# set flag source-active | match 230.0.0.3
   ```

7. View the trace file.

   ```
   user@host> file list /var/log
   user@host> file show /var/log/msdp-trace
   ```

Related Documentation

- Understanding MSDP on page 39
- *Junos OS Tracing and Logging Operations* in the *Junos OS Administration Library for Routing Devices*

Configuring the Interface to Accept Traffic from a Remote Source

You can configure an incoming interface to accept traffic from a remote source. A remote source is a source that is not on the same subnet as the incoming interface. This enables the remote source to be learned and advertised by MSDP so that receivers in other MSDP areas can join the source. You do not need to disable RPF checking, but you do need to ensure that the best path to reach the remote source is through the incoming interface.

In this sample configuration, the incoming interface (**ge-1/3/0**) is on a provider edge (PE) router on the receiver side of a multicast VPN.

To accept traffic from a remote source:

1. Edit the incoming interface.

   ```
   [edit protocols pim interface ge-1/3/0.0]
   user@host# set accept-remote-source
   ```

2. If the incoming interface is not the only way to reach the remote source, ensure that the best path to reach the remote source is through the incoming interface. One way to do this is to use AS path prepending on the other possible routes.

```
[edit policy-options policy-statement as-path-prepend term prepend]
user@host# set from route-filter 192.168.0.0/16 orlonger
user@host# set from route-filter 172.16.0.0/16 orlonger
user@host# set then as-path-prepend "1111"
```

Another way to do this might be to configure a static route on the receiver side PE router to the source.

4. After the configuration is committed, use the **show pim statistics** and **show msdp source** commands to verify that the interface is accepting traffic from the remote source.

Related Documentation

- *Example: Allowing MBGP MVPN Remote Sources*
- *Understanding Prepending AS Numbers to BGP AS Paths* in the *Routing Policy Feature Guide for Routing Devices*
- show msdp source on page 375 in the CLI Explorer
- show pim statistics on page 448 in the CLI Explorer

Example: Configuring MSDP

Configure a router to act as a PIM sparse-mode rendezvous point and an MSDP peer:

```
[edit]
routing-options {
  interface-routes {
    rib-group ifrg;
  }
  rib-groups {
    ifrg {
      import-rib [inet.0 inet.2];
    }
    mcrg {
      export-rib inet.2;
      import-rib inet.2;
    }
  }
}
protocols {
  bgp {
    group lab {
      type internal;
      family any;
      neighbor 192.168.6.18 {
        local-address 192.168.6.17;
      }
    }
  }
  pim {
    dense-groups {
      224.0.1.39/32;
      224.0.1.40/32;
    }
```

```
        rib-group mcrg;
        rp {
          local {
            address 192.168.1.1;
          }
        }
        interface all {
          mode sparse-dense;
          version 1;
        }
      }
      msdp {
        rib-group mcrg;
        group lab {
          peer 192.168.6.18 {
            local-address 192.168.6.17;
          }
        }
      }
    }
```

Example: Configuring MSDP with Active Source Limits and Mesh Groups

This example shows how to configure MSDP to filter source-active messages and limit the flooding of source-active messages.

Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable PIM sparse mode. See "PIM Overview" on page 3.
- Configure the router as a PIM sparse-mode RP. See "Configuring Local PIM RPs" on page 75.

Overview

A router interested in MSDP messages, such as an RP, might have to process a large number of MSDP messages, especially source-active messages, arriving from other routers. Because of the potential need for a router to examine, process, and create state tables for many MSDP packets, there is a possibility of an MSDP-based denial-of-service (DoS) attack on a router running MSDP. To minimize this possibility, you can configure

the router to limit the number of source active messages the router accepts. Also, you can configure a threshold for applying random early discard (RED) to drop some but not all MSDP active source messages. Beginning with Junos OS 12.2, you can optionally configure a warning threshold so the device can log warning messages in the system log when a certain number of source-active messages have been received. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of source-active messages have been received. These log messages convey when the configured message limit has been exceeded, when the configured warning threshold has been exceeded, and when the number of messages drop below the configured warning threshold.

By default, the router accepts 25,000 source active messages before ignoring the rest. The limit can be from 1 through 1,000,000. The limit is applied to both the number of messages and the number of MSDP peers.

By default, the router accepts 24,000 source-active messages before applying the RED profile to prevent a possible DoS attack. This number can also range from 1 through 1,000,000. The next 1000 messages are screened by the RED profile and the accepted messages processed. If you configure no drop profiles (as this example does not), RED is still in effect and functions as the primary mechanism for managing congestion. In the default RED drop profile, when the packet queue fill-level is 0 percent, the drop probability is 0 percent. When the fill-level is 100 percent, the drop probability is 100 percent.

NOTE: The router ignores source-active messages with encapsulated TCP packets. Multicast does not use TCP; segments inside source-active messages are most likely the result of worm activity.

The number configured for the threshold must be less than the number configured for the maximum number of active MSDP sources.

The warning threshold is a percentage of maximum number of MSDP source-active messages received, so you must configure the source-active message limit to configure a warning threshold. The range for the warning threshold is 1 through 100 percent. You can further specify the amount of time (in seconds) between the log messages. The range is 6 through 32,767 seconds.

You can configure an active source limit globally, for a group, or for a peer. If active source limits are configured at multiple levels of the hierarchy (as shown in this example), all are applied.

You can configure an active source limit for an address range as well as for a specific peer. A per-source active source limit uses an IP prefix and prefix length instead of a specific address. You can configure more than one per-source active source limit. The longest match determines the limit.

Per-source active source limits can be combined with active source limits at the peer, group, and global (instance) hierarchy level. Per-source limits are applied before any other type of active source limit. Limits are tested in the following order:

- Per-source

- Per-peer or group
- Per-instance

An active source message must “pass” all limits established before being accepted. For example, if a source is configured with an active source limit of 10,000 active multicast groups and the instance is configured with a limit of 5000 (and there are no other sources or limits configured), only 5000 active source messages are accepted from this source.

MSDP mesh groups are groups of peers configured in a full-mesh topology that limits the flooding of source-active messages to neighboring peers. Every mesh group member must have a peer connection with every other mesh group member. When a source-active message is received from a mesh group member, the source-active message is always accepted but is not flooded to other members of the same mesh group. However, the source-active message is flooded to non-mesh group peers or members of other mesh groups. By default, standard flooding rules apply if **mesh-group** is not specified.

CAUTION: When configuring MSDP mesh groups, you must configure all members the same way. If you do not configure a full mesh, excessive flooding of source-active messages can occur.

A common application for MSDP mesh groups is peer-reverse-path-forwarding (peer-RPF) check bypass. For example, if there are two MSDP peers inside an autonomous system (AS), and only one of them has an external MSDP session to another AS, the internal MSDP peer often rejects incoming source-active messages relayed by the peer with the external link. Rejection occurs because the external MSDP peer must be reachable by the internal MSDP peer through the next hop toward the source in another AS, and this next-hop condition is not certain. To prevent rejections, configure an MSDP mesh group on the internal MSDP peer so it always accepts source-active messages.

NOTE: An alternative way to bypass the peer-RPF check is to configure a default peer. In networks with only one MSDP peer, especially stub networks, the source-active message always needs to be accepted. An MSDP default peer is an MSDP peer from which all source-active messages are accepted without performing the peer-RPF check. You can establish a default peer at the peer or group level by including the default-peer statement.

Table 7 on page 152 explains how flooding is handled by peers in this example. Figure 18 on page 153 illustrates source-active message flooding between different mesh groups and peers within the same mesh group.

Table 7: Source-Active Message Flooding Explanation

Source-Active Message Received From	Source-Active Message Flooded To	Source-Active Message Not Flooded To
Peer 21	Peer 11, Peer 12, Peer 13, Peer 31, Peer 32	Peer 22

Table 7: Source-Active Message Flooding Explanation *(continued)*

Source-Active Message Received From	Source-Active Message Flooded To	Source-Active Message Not Flooded To
Peer 11	Peer 21, Peer 22, Peer 31, Peer 32	Peer 12, Peer 13
Peer 31	Peer 21, Peer 22, Peer 11, Peer 12, Peer 13, Peer 32	–

Figure 18: Source-Active Message Flooding

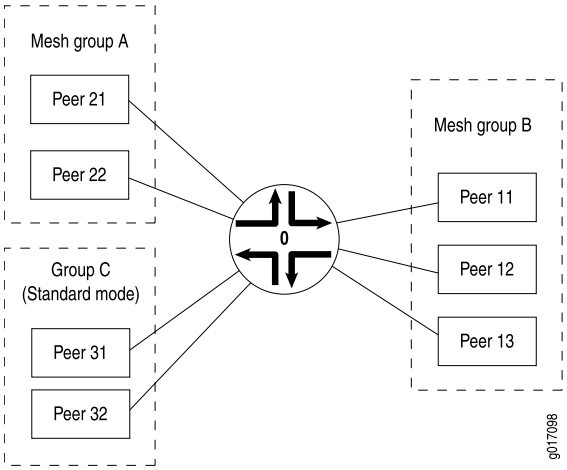

This example includes the following settings:

- **active-source-limit maximum 10000**—Applies a limit of 10,000 active sources to all other peers.
- **active-source-limit log-warning 80**—(Optional) Applies a warning threshold of 80 percent. In this example, the active source maximum is 10,000, so the device will start logging warning messages once it receives 8,000 active source messages.
- **active-source-limit log-interval 20**—(Optional) Applies a 20 second waiting period between system log messages.
- **data-encapsulation disable**—On an RP router using MSDP, disables the default encapsulation of multicast data received in MSDP register messages inside MSDP source-active messages.

  MSDP data encapsulation mainly concerns bursty sources of multicast traffic. Sources that send only one packet every few minutes have trouble with the timeout of state relationships between sources and their multicast groups (S,G). Routers lose data while they attempt to reestablish (S,G) state tables. As a result, multicast register messages contain data, and this data encapsulation in MSDP source-active messages can be turned on or off through configuration.

  By default, MSDP data encapsulation is enabled. An RP running MSDP takes the data packets arriving in the source's register message and encapsulates the data inside an MSDP source-active message.

However, data encapsulation creates both a multicast forwarding cache entry in the **inet.1** table (this is also the forwarding table) and a routing table entry in the **inet.4** table. Without data encapsulation, MSDP creates only a routing table entry in the **inet.4** table. In some circumstances, such as the presence of Internet worms or other forms of DoS attack, the router's forwarding table might fill up with these entries. To prevent the forwarding table from filling up with MSDP entries, you can configure the router not to use MSDP data encapsulation. However, if you disable data encapsulation, the router ignores and discards the encapsulated data. Without data encapsulation, multicast applications with bursty sources having transmit intervals greater than about 3 minutes might not work well.

- **group MSDP-group local-address 10.1.2.3**—Specifies the address of the local router (this router).
- **group MSDP-group mode mesh-group**—Specifies that all peers belonging to the **MSDP-group** group are mesh group members.
- **group MSDP-group peer 10.10.10.10**—Prevents the sending of source-active messages to neighboring peer 10.10.10.10.
- **group MSDP-group peer 10.10.10.10 active-source-limit maximum 7500**—Applies a limit of 7500 active sources to MSDP peer 10.10.10.10 in group **MSDP-group**.
- **peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000**—Applies a threshhold of 4000 active sources and a limit of 5000 active sources to MSDP peer 10.0.0.1.
- **source 10.1.0.0/16 active-source-limit maximum 500**—Applies a limit of 500 active sources to any source on the 10.1.0.0/16 network.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols msdp data-encapsulation disable
set protocols msdp active-source-limit maximum 10000
set protocols msdp active-source-limit log-warning 80
set protocols msdp active-source-limit log-interval 20
set protocols msdp peer 10.0.0.1 active-source-limit maximum 5000
set protocols msdp peer 10.0.0.1 active-source-limit threshold 4000
set protocols msdp source 10.1.0.0/16 active-source-limit maximum 500
set protocols msdp group MSDP-group mode mesh-group
set protocols msdp group MSDP-group local-address 10.1.2.3
set protocols msdp group MSDP-group peer 10.10.10.10 active-source-limit maximum
  7500
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure MSDP source active routes and mesh groups:

1. (Optional) Disable data encapsulation.

```
[edit protocols msdp]
user@host# set data-encapsulation disable
```

2. Configure the active source limits.

```
[edit protocols msdp]
user@host# set peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000
user@host# set group MSDP-group peer 10.10.10.10 active-source-limit maximum
  7500
user@host# set active-source-limit maximum 10000
user@host# set source 10.1.0.0/16 active-source-limit maximum 500
```

3. (Optional) Configure the threshold at which warning messages are logged and the amount of time between log messages.

```
[edit protocols msdp]
user@host# set active-source-limit log-warning 80
user@host# set active-source-limit log-interval 20
```

4. Configure the mesh group.

```
[edit protocols msdp]
user@host# set group MSDP-group mode mesh-group
user@host# set group MSDP-group peer 10.10.10.10
user@host# set group MSDP-group local-address 10.1.2.3
```

5. If you are done configuring the device, commit the configuration.

```
[edit routing-instances]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols** command.

```
user@host# show protocols
msdp {
  data-encapsulation disable;
  active-source-limit {
    maximum 10000;
    log-warning 80;
    log-interval 20;
  }
  peer 10.0.0.1 {
    active-source-limit {
      maximum 5000;
      threshold 4000;
    }
  }
```

```
    source 10.1.0.0/16 {
      active-source-limit {
        maximum 500;
      }
    }
    group MSDP-group {
      mode mesh-group;
      local-address 10.1.2.3;
      peer 10.10.10.10 {
        active-source-limit {
          maximum 7500;
        }
      }
    }
  }
```

Verification

To verify the configuration, run the following commands:

- show msdp source-active
- show msdp statistics

Related Documentation

- *Example: Configuring MSDP in a Routing Instance*
- Filtering MSDP SA Messages on page 40
- *Configuring RED Drop Profiles* in the *Junos OS Class of Service Library for Routing Devices*
- Configuring Local PIM RPs on page 75

Example: Configuring PIM Anycast With or Without MSDP

When you configure anycast RP, you bypass the restriction of having one active rendezvous point (RP) per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP stops operating, sources and receivers are taken to a new RP by means of unicast routing.

You can configure anycast RP to use PIM and MSDP for IPv4, or PIM alone for both IPv4 and IPv6 scenarios. Both are discussed in this section.

We recommend a static RP mapping with anycast RP over a bootstrap router and auto-RP configuration because it provides all the benefits of a bootstrap router and auto-RP without the complexity of the BSR and auto-RP mechanisms.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default RP mode (at the **[edit protocols pim rp static address *address*]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface *interface-name*]** hierarchy level). Explicitly configured versions override the defaults. This example explicitly configures PIMv2 on the interfaces.

The following example shows an anycast RP configuration for the RP routers, first with MSDP and then using PIM alone, and for non-RP routers.

1. For a network using an RP with MSDP, configure the RP using the **lo0** loopback interface, which is always up. Include the **address** statement and specify the unique and routable router ID and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement for the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

   ```
   interfaces {
     lo0 {
       description "PIM RP";
       unit 0 {
         family inet {
           address 198.58.3.254/32;
           primary;
           address 198.58.3.253/32;
         }
       }
     }
   }
   ```

2. Specify the RP address. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level (the same address as the secondary **lo0** interface).

   For all interfaces, include the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

   ```
   protocols {
     pim {
       rp {
         local {
           family inet;
           address 198.58.3.253;
         }
         interface all {
           mode sparse;
           version 2;
         }
         interface fxp0.0 {
           disable;
         }
       }
     }
   }
   ```

3. Configure MSDP peering. Include the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, include the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
  msdp {
    peer 198.58.3.250 {
      local-address address 198.58.3.254;
    }
  }
}
```

NOTE: If you need to configure a PIM RP for both IPv4 and IPv6 scenarios, perform Step 4 and Step 5. Otherwise, go to Step 6.

4. Configure an RP using the **lo0** loopback interface, which is always up. Include the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement on the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32 {
          primary;
        }
        address 198.58.3.253/32;
      }
    }
  }
}
```

5. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

   For all interfaces, include the **mode** statement to set the mode to **sparse**, and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by Including the **disable** statement for that interface.

   Include the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

```
protocols {
  pim {
    rp {
      local {
        family inet {
          address 198.58.3.253;
          anycast-pim {
            rp-set {
              address 198.58.3.240;
              address 198.58.3.241 forward-msdp-sa;
            }
            local-address 198.58.3.254; #If not configured, use lo0 primary
          }
        }
      }
    }
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at the **[edit protocols pim rp static]** hierarchy level. Include the **version** statement at the **[edit protocols pim rp static address]** hierarchy level to specify PIM version 2.

```
protocols {
  pim {
    rp {
      static {
        address 198.58.3.253 {
          version 2;
        }
      }
    }
  }
}
```

7. Include the **mode** statement at the **[edit protocols pim interface all]** hierarchy level to specify sparse mode on all interfaces. Then include the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
```

```
        }
        interface fxp0.0 {
          disable;
        }
      }
    }
```

Configuring a PIM Anycast RP Router with MSDP

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
protocols {
  pim {
    rp {
      local {
        family inet;
        address 198.58.3.253;
      }
      interface all {
        mode sparse;
        version 2;
      }
      interface fxp0.0 {
        disable;
      }
    }
  }
}
```

To configure MSDP peering, add the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, add the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
  msdp {
    peer 198.58.3.250 {
      local-address 198.58.3.254;
    }
  }
}
```

CHAPTER 26

Source-Specific Multicast

Example: Configuring PIM SSM on a Network

The following example shows how PIM SSM is configured between a receiver and a source in the network illustrated in Figure 19 on page 161.

Figure 19: Network on Which to Configure PIM SSM

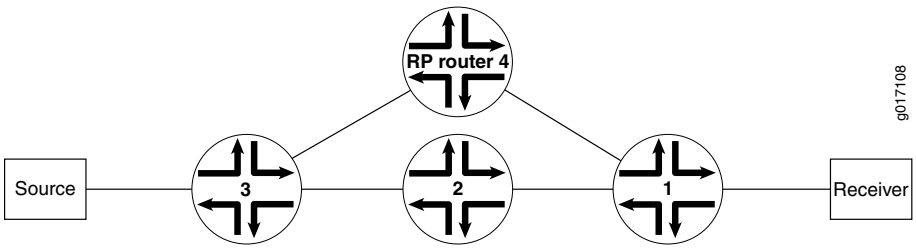

This example shows how to configure the IGMP version to IGMPv3 on all receiving host interfaces.

1. Enable IGMPv3 on all host-facing interfaces, and disable IGMP on the **fxp0.0** interface on Router 1.

   user@router1# set protocols igmp interface all version 3
   user@router1# set protocols igmp interface fxp0.0 disable

   NOTE: **When you configure IGMPv3 on a router, hosts on interfaces configured with IGMPv2 cannot join the source tree.**

2. After the configuration is committed, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@router1> show configuration protocol igmp

[edit protocols igmp]
interface all {
  version 3;
}
interface fxp0.0 {
  disable;
}
```

3. Use the **show igmp interface** command to verify that IGMP interfaces are configured.

```
user@router1> show igmp interface
Interface      State    Querier      Timeout Version Groups
fe-0/0/0.0     Up       198.58.3.245     213       3      0
fe-0/0/1.0     Up       198.58.3.241     220       3      0
fe-0/0/2.0     Up       198.58.3.237     218       3      0
Configured Parameters:
IGMP Query Interval (1/10 secs): 1250
IGMP Query Response Interval (1/10 secs): 100
IGMP Last Member Query Interval (1/10 secs): 10
IGMP Robustness Count: 2
Derived Parameters:
IGMP Membership Timeout (1/10 secs): 2600
IGMP Other Querier Present Timeout (1/10 secs): 2550
```

4. Use the **show pim join extensive** command to verify the PIM join state on Router 2 and Router 3 (the upstream routers).

```
user@router2> show pim join extensive
232.1.1.1       10.4.1.2                        sparse
     Upstream interface: fe-1/1/3.0
     Upstream State: Local Source
     Keepalive timeout: 209
     Downstream Neighbors:
        Interface: so-1/0/2.0
            10.10.71.1       State: Join   Flags: S    Timeout: 209
```

5. Use the **show pim join extensive** command to verify the PIM join state on Router 1 (the router connected to the receiver).

```
user@router1> show pim join extensive
232.1.1.1       10.4.1.2                        sparse
     Upstream interface: so-1/0/2.0
     Upstream State: Join to Source
     Keepalive timeout: 209
     Downstream Neighbors:
        Interface: fe-0/2/3.0
            10.3.1.1         State: Join   Flags: S    Timeout: Infinity
```

Example: Configuring an SSM-Only Domain

Deploying an SSM-only domain is much simpler than deploying an ASM domain because it only requires a few configuration steps. Enable PIM sparse mode on all interfaces by adding the **mode** statement at the **[edit protocols pim interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface. Then configure IGMPv3 on all host-facing interfaces by adding the **version** statement at the **[edit protocols igmp interface *interface-name*]** hierarchy level.

In the following example, the host-facing interface is **fe-0/1/2**:

```
[edit]
protocols {
   pim {
      interface all {
         mode sparse;
         version 2;
      }
      interface fxp0.0 {
         disable;
      }
   }
   igmp {
      interface fe-0/1/2 {
         version 3;
      }
   }
}
```

Example: Configuring SSM Mapping

SSM mapping does not require that all hosts support IGMPv3. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report. This enables hosts running IGMPv1 or IGMPv2 to participate in SSM until the hosts transition to IGMPv3.

SSM mapping applies to all group addresses that match the policy, not just those that conform to SSM addressing conventions (232/8 for IPv4, ff30::/32 through ff3F::/32 for IPv6).

We recommend separate SSM maps for IPv4 and IPv6 if both address families require SSM support. If you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv4 context (using IGMP), only the IPv4 addresses in the list are used. If there are no such addresses, no action is taken. Similarly, if you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv6 context (using MLD), only the IPv6 addresses in the list are used. If there are no such addresses, no action is taken.

In this example, you create a policy to match the group addresses that you want to translate to IGMPv3. Then you define the SSM map that associates the policy with the source addresses where these group addresses are found. Finally, you apply the SSM map to one or more IGMP (for IPv4) or MLD (for IPv6) interfaces.

1. Create an SSM policy named **ssm-policy-example**. The policy terms match the IPv4 SSM group address 232.1.1.1/32 and the IPv6 SSM group address ff35::1/128. All other addresses are rejected.

   ```
   user@router1# set policy-options policy-statement ssm-policy-example term A from
     route-filter 232.1.1.1/32 exact
   user@router1# set policy-options policy-statement ssm-policy-example term A then
     accept
   user@router1# set policy-options policy-statement ssm-policy-example term B from
     route-filter ff35::1/128 exact
   ```

```
user@router1# set policy-options policy-statement ssm-policy-example term B then
  accept
```

2. After the configuration is committed, use the **show configuration policy-options** command to verify the policy configuration.

```
user@host> show configuration policy-options
```

```
[edit policy-options]
policy-statement ssm-policy-example {
  term A {
    from {
      route-filter 232.1.1.1/32 exact;
    }
    then accept;
  }
  term B {
    from {
      route-filter ff35::1/128 exact;
    }
    then accept;
  }
  then reject;
}
```

The group addresses must match the configured policy for SSM mapping to occur.

3. Define two SSM maps, one called **ssm-map-ipv6-example** and one called **ssm-map-ipv4-example**, by applying the policy and configuring the source addresses as a multicast routing option.

```
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example policy
  ssm-policy-example
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example source
  fec0::1 fec0::12
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example policy
  ssm-policy-example
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
  10.10.10.4
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
  192.168.43.66
```

4. After the configuration is committed, use the **show configuration routing-options** command to verify the policy configuration.

```
user@host> show configuration routing-options
```

```
[edit routing-options]
multicast {
  ssm-map ssm-map-ipv6-example {
    policy ssm-policy-example;
    source [ fec0::1 fec0::12 ];
  }
  ssm-map ssm-map-ipv4-example {
    policy ssm-policy-example;
    source [ 10.10.10.4 192.168.43.66 ];
  }
}
```

We recommend separate SSM maps for IPv4 and IPv6.

5. Apply SSM maps for IPv4-to-IGMP interfaces and SSM maps for IPv6-to-MLD interfaces:

   user@host# set protocols igmp interface fe-0/1/0.0 ssm-map ssm-map-ipv4-example
   user@host# set protocols mld interface fe-0/1/1.0 ssm-map ssm-map-ipv6-example

6. After the configuration is committed, use the **show configuration protocol** command to verify the IGMP and MLD protocol configuration.

   user@router1> show configuration protocol

   ```
   [edit protocols]
   igmp {
     interface fe-0/1/0.0 {
       ssm-map ssm-map-ipv4-example;
     }
   }
   mld {
     interface fe-/0/1/1.0 {
       ssm-map ssm-map-ipv6-example;
     }
   }
   ```

7. Use the **show igmp interface** and the **show mld interface** commands to verify that the SSM maps are applied to the interfaces.

   ```
   user@host> show igmp interface fe-0/1/0.0
   Interface: fe-0/1/0.0
        Querier: 192.168.224.28
        State:         Up Timeout:    None Version:   2 Groups:   2
        SSM Map: ssm-map-ipv4-example

   user@host> show mld interface fe-0/1/1.0
   Interface: fe-0/1/1.0
        Querier: fec0:0:0:0:1::12
        State:         Up Timeout:    None Version:   2 Groups:   2
        SSM Map: ssm-map-ipv6-example
   ```

Example: Configuring Source-Specific Multicast Groups with Any-Source Override

This example shows how to extend source-specific multicast (SSM) group operations beyond the default IP address range of 232.0.0.0 through 232.255.255.255. This example also shows how to accept any-source multicast (ASM) join messages (*,G) for group addresses that are within the default or configured range of SSM groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

Requirements

Before you begin, configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.

Overview

To deploy SSM, configure PIM sparse mode on all routing device interfaces and issue the necessary SSM commands, including specifying IGMPv3 or MLDv2 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group members interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3 and MLDv2, are used in PIM SSM. Only sources that are specified send traffic to the SSM group.

In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3 or MLDv2) to join group G and source S (see Figure 20 on page 166). The directly connected PIM sparse-mode router, the receiver's designated router (DR), sends an (S,G) join message to its reverse-path forwarding (RPF) neighbor for the source. Notice in Figure 20 on page 166 that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 20: Receiver Sends Messages to Join Group G and Source S

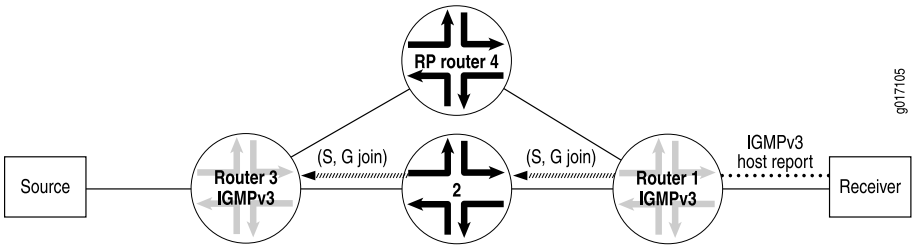

The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In Figure 21 on page 166, the source tree is built across the network to Router 3, the last-hop router connected to the source.

Figure 21: Router 3 (Last-Hop Router) Joins the Source Tree

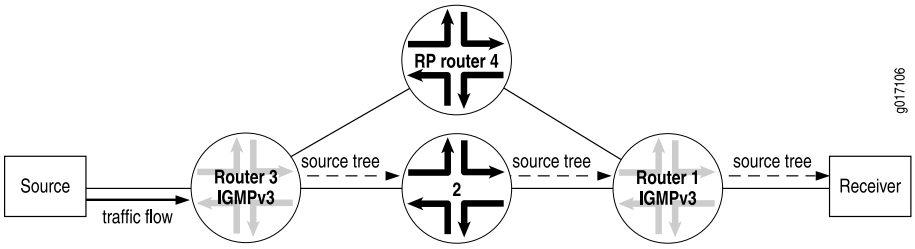

Using the source tree, multicast traffic is delivered to the subscribing host (see Figure 22 on page 167).

Figure 22: (S,G) State Is Built Between the Source and the Receiver

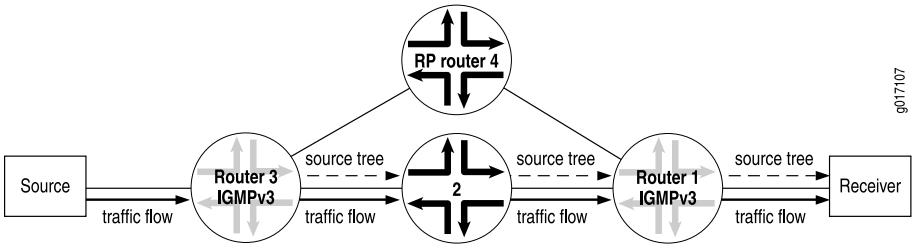

SSM can operate in include mode or in exclude mode. In exclude mode the receiver specifies a list of sources that it does not want to receive the multicast group traffic from. The routing device forwards traffic to the receiver from any source except the sources specified in the exclusion list. The receiver accepts traffic from any sources except the sources specified in the exclusion list.

This example works with the simple RPF topology shown in Figure 23 on page 167.

Figure 23: Simple RPF Topology

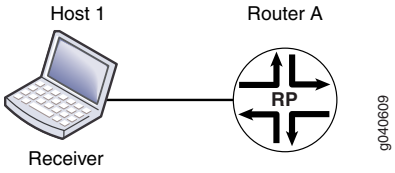

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 10.255.72.46
set protocols pim rp local group-ranges 239.0.0.0/24
set protocols pim interface fe-1/0/0.0 mode sparse
set protocols pim interface lo0.0 mode sparse
set routing-options multicast ssm-groups 232.0.0.0/8
set routing-options multicast ssm-groups 239.0.0.0/8
set routing-options multicast asm-override-ssm
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an RPF policy:

1. Configure OSPF.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface fxp0.0 disable
user@host# set area 0.0.0.0 interface all
```

2. Configure PIM sparse mode.

```
[edit protocols pim]
user@host# set rp local address 10.255.72.46
user@host# set rp local group-ranges 239.0.0.0/24
user@host# set interface fe-1/0/0.0 mode sparse
user@host# set interface lo0.0 mode sparse
```

3. Configure additional SSM groups.

```
[edit routing-options]
user@host# set ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ]
```

4. Configure the RP to accept ASM join messages for groups within the SSM address range.

```
[edit routing-options]
user@host# set multicast asm-override-ssm
```

5. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols** and **show routing-options** commands.

```
user@host# show protocols
ospf {
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface all;
    }
}
pim {
    rp {
        local {
            address 10.255.72.46;
            group-ranges {
                239.0.0.0/24;
            }
        }
    }
    interface fe-1/0/0.0 {
        mode sparse;
    }
    interface lo0.0 {
        mode sparse;
    }
}
```

```
user@host# show routing-options
multicast {
    ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ];
```

```
    asm-override-ssm;
}
```

Verification

To verify the configuration, run the following commands:

- show igmp group
- show igmp statistics
- show pim join

Related Documentation

- Source-Specific Multicast Groups Overview on page 43

Example: Configuring SSM Maps for Different Groups to Different Sources

Multiple SSM Maps and Groups for Interfaces

You can configure multiple source-specific multicast (SSM) maps so that different groups map to different sources, which enables a single multicast group to map to different sources for different interfaces.

Example: Configuring Multiple SSM Maps Per Interface

This example shows how to assign more than one SSM map to an IGMP interface.

Requirements

This example requires Junos OS Release 11.4 or later.

Overview

In this example, you configure a routing policy, POLICY-ipv4-example1, that maps multicast group join messages over an IGMP logical interface to IPv4 multicast source addresses based on destination IP address as follows:

Routing Policy Name	Multicast Group Join Messages for a Route Filter at This Destination Address	Multicast Source Addresses
POLICY-ipv4-example1 term 1	232.1.1.1	10.10.10.4, 192.168.43.66

Routing Policy Name	Multicast Group Join Messages for a Route Filter at This Destination Address	Multicast Source Addresses
POLICY-ipv4-example1 term 2	232.1.1.2	10.10.10.5, 192.168.43.67

You apply routing policy POLICY-ipv4-example1 to IGMP logical interface fe-0/1/0.0.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure this example, perform the following task:

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set policy-options policy-statement POLICY-ipv4-example1 term 1 from route-filter
  232.1.1.1/32 exact
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source
  10.10.10.4
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source
  192.168.43.66
set policy-options policy-statement POLICY-ipv4-example1 term 1 then accept
set policy-options policy-statement POLICY-ipv4-example1 term 2 from route-filter
  232.1.1.2/32 exact
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source
  10.10.10.5
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source
  192.168.43.67
set policy-options policy-statement POLICY-ipv4-example1 term 2 then accept
set protocols igmp interface fe-0/1/0.0 ssm-map-policy POLICY-ipv4-example1
```

Step-by-Step Procedure

To configure multiple SSM maps per interface:

1. Configure protocol-independent routing options for route filter 232.1.1.1, and specify the multicast source addresses to which matching multicast groups are to be mapped.

   ```
   [edit policy-options policy-statement POLICY-ipv4-example1 term 1]
   user@host# set from route-filter 232.1.1.1/32 exact
   user@host# set then ssm-source 10.10.10.4
   user@host# set then ssm-source 192.168.43.66
   user@host# set then accept
   ```

2. Configure protocol-independent routing options for route filter 232.1.1.2, and specify the multicast source addresses to which matching multicast groups are to be mapped.

   ```
   [edit policy-options policy-statement POLICY-ipv4-example1 term 2]
   user@host# set from route-filter 232.1.1.2/32 exact
   user@host# set then ssm-source 10.10.10.5
   ```

```
user@host# set then ssm-source 192.168.43.67
user@host# set then accept
```

3. Apply the policy map POLICY-ipv4-example1 to IGMP logical interface fe-0/1/1/0.

```
[edit protocols igmp interface fe-0/1/0.0]
user@host# set ssm-map-policy POLICY-ipv4-example1
```

Results After the configuration is committed, confirm the configuration by entering the **show policy-options** and **show protocols** configuration mode commands. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
user@host#> show policy-options
policy-statement POLICY-ipv4-example1 {
  term 1 {
    from {
      route-filter 232.1.1.1/32 exact;
    }
    then {
      ssm-source [ 10.10.10.4 192.168.43.66 ];
      accept;
    }
  }
  term 2{
    from {
      route-filter 232.1.1.2/32 exact;
    }
    then {
      ssm-source [ 10.10.10.5 192.168.43.67 ];
      accept;
    }
  }
}

user@host# show protocols
igmp {
  interface fe-0/1/0.0 {
    ssm-map-policy POLICY-ipv4-example1;
  }
}
```

Verification

Confirm that the configuration is working properly.

- Displaying Information About IGMP-Enabled Interfaces on page 171
- Displaying the PIM Groups on page 172
- Displaying the Entries in the IP Multicast Forwarding Table on page 172

Displaying Information About IGMP-Enabled Interfaces

Purpose Verify that the SSM map policy POLICY-ipv4-example1 is applied to logical interface fe-0/1/0.0.

Action Use the show igmp interface operational mode command for the IGMP logical interface to which you applied the SSM map policy.

```
user@host> show igmp interface
Interface: fe-0/1/0.0
    Querier: 10.111.30.1
    State:         Up Timeout:    None Version:  2 Groups:      2
    SSM Map Policy: POLICY-ipv4-example1;

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

The command output displays the name of IGMP logical interface (fe-0/1/0.0), the address of the routing device that has been elected to send membership queries and group information.

Displaying the PIM Groups

Purpose Verify the Protocol Independent Multicast (PIM) source and group pair (S,G) entries.

Action Use the show pim join **extensive 232.1.1.1** operational mode command to display the PIM source and group pair (S,G) entries for the 232.1.1.1 group.

Displaying the Entries in the IP Multicast Forwarding Table

Purpose Verify that the IP multicast forwarding table displays the mroute state.

Action Use the show multicast route **extensive** operational mode command to display the entries in the IP multicast forwarding table to verify that the **Route state** is active and that the **Forwarding state** is forwarding.

Related Documentation

- *Example: Configuring Source-Specific Multicast*
- *Example: Configuring Source-Specific Draft-Rosen 7 Multicast VPNs*

CHAPTER 27

PIM Configuration Statements

- version (BFD) on page 241
- version (PIM) on page 242
- wildcard-source (PIM RPF Selection) on page 243

address (Anycast RPs)

Syntax address *address* <forward-msdp-sa>;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp local (inet | inet6) anycast-pim rp-set],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local (inet | inet6) anycast-pim rp-set],
[edit protocols pim rp local (inet | inet6) anycast-pim rp-set],
[edit routing-instances *routing-instance-name* protocols pim rp local (inet | inet6) anycast-pim rp-set]

Release Information Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the anycast rendezvous point (RP) addresses in the RP set. Multiple addresses can be configured in an RP set. If the RP has peer Multicast Source Discovery Protocol (MSDP) connections, then the RP must forward MSDP source active (SA) messages.

Options ***address***—RP address in an RP set.

forward-msdp-sa—(Optional) Forward MSDP SAs to this address.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

address (Local RPs)

Syntax address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the local rendezvous point (RP) address.

Options ***address***—Local RP address.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Local PIM RPs on page 75

address (Static RPs)

Syntax

```
address address {
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  override;
  version version;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp static],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp static],
[edit protocols pim static],
[edit routing-instances *routing-instance-name* protocols pim rp static]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure static rendezvous point (RP) addresses. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.

For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.

Options

address—Static RP address.

Default: 224.0.0.0/4

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Static PIM RP Address on the Non-RP Routing Device on page 77

algorithm

Syntax algorithm *algorithm-name*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection authentication],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection authentication]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the algorithm to use for BFD authentication.

Options ***algorithm-name***—Name of algorithm to use for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text. One or more passwords can be configured.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive rates greater than 100 ms.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive rates greater than 100 ms.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding Bidirectional Forwarding Detection Authentication for PIM*
- Configuring BFD Authentication for PIM on page 108
- authentication on page 181

anycast-pim

Syntax

```
anycast-pim {
  rp-set {
    address address <forward-msdp-sa>;
  }
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information

Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure properties for anycast RP using PIM.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PIM Anycast With or Without MSDP on page 79

assert-timeout

Syntax assert-timeout *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Multicast routing devices running PIM sparse mode often forward the same stream of multicast packets onto the same LAN through the rendezvous-point tree (RPT) and shortest-path tree (SPT). PIM assert messages help routing devices determine which routing device forwards the traffic and prunes the RPT for this group. By default, routing devices enter an assert cycle every 180 seconds. You can configure this assert timeout to be between 5 and 210 seconds.

Options ***seconds***—Time for routing device to wait before another assert message cycle.

Range: 5 through 210 seconds

Default: 180 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring the PIM Assert Timeout on page 99

authentication (Protocols PIM)

Syntax

```
authentication {
  algorithm algorithm-name;
  key-chain key-chain-name;
  loose-check;
}
```

Hierarchy Level

[edit protocols pim interface *interface-name* family (inet | inet6) bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface family (inet | inet6) *interface-name* bfd-liveness-detection]

Release Information

Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure the algorithm, security keychain, and level of authentication for BFD sessions running on PIM interfaces.

The remaining statements are explained separately.

Options

The statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD Authentication for PIM on page 108
- Configuring BFD for PIM on page 107
- *Understanding Bidirectional Forwarding Detection Authentication for PIM*
- bfd-liveness-detection on page 182
- key-chain (Protocols PIM) on page 205
- loose-check on page 208

bfd-liveness-detection (Protocols PIM)

Syntax

```
bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (0 | 1 | automatic);
}
```

Hierarchy Level

[edit protocols pim interface *interface-name* family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* family (inet | inet6)]

Release Information

Statement introduced in Junos OS Release 8.1.
authentication option introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure bidirectional forwarding detection (BFD) timers and authentication for PIM.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 107
- Configuring BFD Authentication for PIM on page 108

bootstrap

Syntax

```
bootstrap {
  family (inet | inet6) {
    export [ policy-names ];
    import [ policy-names ];
    priority number;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure parameters to control bootstrap routers and messages.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring PIM Bootstrap Properties for IPv4*
- *Configuring PIM Bootstrap Properties for IPv4 or IPv6*

bootstrap-export

Syntax bootstrap-export [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more export policies to control outgoing PIM bootstrap messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring PIM Bootstrap Properties for IPv4*
- *Configuring PIM Bootstrap Properties for IPv4 or IPv6*
- bootstrap-import on page 185

bootstrap-import

Syntax bootstrap-import [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more import policies to control incoming PIM bootstrap messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring PIM Bootstrap Properties for IPv4*
- *Configuring PIM Bootstrap Properties for IPv4 or IPv6*
- bootstrap-export on page 184

bootstrap-priority

Syntax bootstrap-priority *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure whether this routing device is eligible to be a bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router.

Options *number*—Priority for becoming the bootstrap router. A value of 0 means that the routing device is not eligible to be the bootstrap router.

Range: 0 through 255

Default: 0

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring PIM Bootstrap Properties for IPv4*

detection-time (BFD for PIM)

Syntax

```
detection-time {
  threshold milliseconds;
}
```

Hierarchy Level

[edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for BFD authentication introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 107
- bfd-liveness-detection on page 182
- threshold on page 235

disable (PIM)

Syntax disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* protocols pim family (inet | inet6)],
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit protocols pim],
[edit protocols pim family (inet | inet6)],
[edit protocols pim interface *interface-name*],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim],
[edit routing-instances *routing-instance-name* protocols pim family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information
Statement introduced before Junos OS Release 7.4.
disable statement extended to the **[family]** hierarchy level in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Explicitly disable PIM at the protocol, interface or family hierarchy levels.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling PIM on page 59
- *disable (PIM Graceful Restart)*

dr-election-on-p2p

Syntax dr-election-on-p2p;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information
Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Enable PIM designated router (DR) election on point-to-point (P2P) links.

Default No PIM DR election is performed on point-to-point links.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Designated Router Election on Point-to-Point Links on page 64

dr-register-policy

Syntax dr-register-policy [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information
Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more policies to control outgoing PIM register messages.

Options *policy-names*—Name of one or more import policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Register Message Filters on a PIM RP and DR on page 96
- rp-register-policy on page 229

embedded-rp

Syntax

```
embedded-rp {
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  maximum-rps limit;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure properties for embedded IP version 6 (IPv6) RPs.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring PIM Embedded RP for IPv6*

export (Protocols PIM Bootstrap)

Syntax export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp bootstrap family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap family (inet | inet6)],
[edit protocols pim rp bootstrap family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap family (inet | inet6)]

Release Information Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more export policies to control outgoing PIM bootstrap messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Configuring PIM Bootstrap Properties for IPv4*
- *Configuring PIM Bootstrap Properties for IPv4 or IPv6*
- import (Protocols PIM Bootstrap) on page 199

export (Protocols PIM)

Syntax export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more export policies to control outgoing PIM join and prune messages. PIM join and prune filters can be applied to PIM-SM and PIM-SSM messages. PIM join and prune filters cannot be applied to PIM-DM messages.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation
- Filtering Outgoing PIM Join Messages on page 94

family (Bootstrap)

Syntax

```
family (inet | inet6) {
    export [ policy-names ];
    import [ policy-names ];
    priority number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp bootstrap],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap],
[edit protocols pim rp bootstrap],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap]

Release Information

Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure which IP protocol type bootstrap properties to apply.

Options

inet—Apply IP version 4 (IPv4) local RP properties.

inet6—Apply IPv6 local RP properties.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring PIM Bootstrap Properties for IPv4*
- *Configuring PIM Bootstrap Properties for IPv4 or IPv6*

family (Protocols PIM)

Syntax

```
family (inet | inet6) {
  disable;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS 11.3 for the QFX Series.

Description

Enable the PIM protocol for the specified family.

Options

inet—Enable the PIM protocol for the IP version 4 (IPv4) address family.

inet6—Enable the PIM protocol for the IP version 6 (IPv6) address family.

The remaining statement is explained separately.

Related Documentation

- Disabling PIM on page 59
- *disable (PIM Graceful Restart)*
- **disable (PIM) on page 188**

family (Local RP)

Syntax

```
family (inet | inet6) {
  disable;
  address address;
  anycast-pim {
    local-address address;
    rp-set {
      address address <forward-msdp-sa>;
    }
  }
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  hold-time seconds;
  override;
  priority number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp local],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local],
[edit protocols pim rp local],
[edit routing-instances *routing-instance-name* protocols pim rp local]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure which IP protocol type local RP properties to apply.

Options

inet—Apply IP version 4 (IPv4) local RP properties.

inet6—Apply IPv6 local RP properties.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs on page 75

group (RPF Selection)

Syntax

```
group group-address{
  sourcesource-address{
    next-hop next-hop-address;
  }
  wildcard-source {
    next-hop next-hop-address;
  }
}
```

Hierarchy Level [edit routing-instances *routing-instance-name* edit protocols pim rpf-selection]

Release Information Statement introduced in JUNOS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the PIM group address for which you configure RPF selectiongroup (RPF Selection).

Default By default, PIM RPF selection is not configured.

Options ***group-address***—PIM group address for which you configure RPF selection.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- *Example: Configuring PIM RPF Selection*

group-ranges

Syntax

```
group-ranges {
  destination-ip-prefix</prefix-length>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp bidirectional address *address*],
[edit logical-systems *logical-system-name* protocols pim rp embedded-rp],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp embedded-rp],
[edit protocols pim rp bidirectional address *address*],
[edit protocols pim rp embedded-rp],
[edit protocols pim rp local family (inet | inet6)],
[edit protocols pim rp static address *address*],
[edit routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit routing-instances *routing-instance-name* protocols pim rp embedded-rp],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp static address *address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional RP addresses introduced in Junos OS Release 12.1.

Description

Configure the address ranges of the multicast groups for which this routing device can be a rendezvous point (RP).

Default

The routing device is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12).

Options

destination-ip-prefix</prefix-length>—Addresses or address ranges for which this routing device can be an RP.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs on page 75 in the *Multicast Protocols Feature Guide for Routing Devices*
- *Configuring PIM Embedded RP for IPv6* in the *Multicast Protocols Feature Guide for Routing Devices*
- *Example: Configuring Bidirectional PIM*

hello-interval (Protocols PIM)

Syntax hello-interval *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Specify how often the routing device sends PIM hello packets out of an interface.

Options ***seconds***—Length of time between PIM hello packets.

Range: 0 through 255

Default: 30 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- hold-time on page 198
- Modifying the PIM Hello Interval on page 55

hold-time (Protocols PIM)

Syntax hold-time *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp bidirectional address *address*],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit protocols pim rp bidirectional address *address*],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional RP addresses introduced in Junos OS Release 12.1.

Description Specify the time period for which a neighbor is to consider the sending routing device (this routing device) to be operative (up).

Options ***seconds***—Hold time.

Range: 0 through 255

Default: 150 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs on page 75 in the *Multicast Protocols Feature Guide for Routing Devices*
- *Example: Configuring Bidirectional PIM*

import (Protocols PIM Bootstrap)

Syntax import [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp bootstrap (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)],
[edit protocols pim rp bootstrap (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)]

Release Information
Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more import policies to control incoming PIM bootstrap messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Configuring PIM Bootstrap Properties for IPv4*
- *Configuring PIM Bootstrap Properties for IPv4 or IPv6*
- export (Protocols PIM Bootstrap) on page 191

import (Protocols PIM)

Syntax import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more policies to routes being imported into the routing table from PIM. Use the **import** statement to filter PIM join messages and prevent them from entering the network.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Filtering Incoming PIM Join Messages on page 95

infinity

Syntax infinity [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim spt-threshold],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim spt-threshold],
[edit protocols pim spt-threshold],
[edit routing-instances *routing-instance-name* protocols pim spt-threshold]

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the **infinity** statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring the PIM SPT Threshold Policy on page 101

interface

Syntax

```
interface (all | interface-name) {
  disable;
  family (inet | inet6) {
    disable;
  }
  hello-interval seconds;
  mode (dense | sparse | sparse-dense);
  neighbor-policy [ policy-names ];
  override-interval milliseconds;
  priority number;
  propagation-delay milliseconds;
  reset-tracking-bit;
  version version;
}
```

Hierarchy Level

[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Enable PIM on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- PIM on Aggregated Interfaces on page 6

join-load-balance

Syntax

```
join-load-balance {
  automatic;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Enable load balancing of PIM join messages across interfaces and routing devices.

Options

automatic—Enables automatic load balancing of PIM join messages. When a new interface or neighbor is introduced into the network, ECMP joins are redistributed with minimal disruption to traffic.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring PIM Make-Before-Break Join Load Balancing*
- Configuring PIM Join Load Balancing on page 66
- *clear pim join-distribution* in the CLI Explorer

join-prune-timeout

Syntax join-prune-timeout *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the timeout for the join state. If the periodic join refresh message is not received before the timeout expires, the join state is removed.

Options **seconds**—Number of seconds to wait for the periodic join message to arrive.

Range: 210 through 240 seconds

Default: 210 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the Join State Timeout on page 69

key-chain (Protocols PIM)

Syntax key-chain *key-chain-name*;

Hierarchy Level [edit protocols pim interface *interface-name* family {inet | inet6} bfd-liveness-detection authentication],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* family {inet | inet6} bfd-liveness-detection authentication]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement modified in Junos OS Release 12.2 to include **family** in the hierarchy level.

Description Specify the security keychain to use for BFD authentication.

Options ***key-chain-name***—Name of the security keychain to use for BFD authentication. The name is a unique integer between **0** and **63**. This must match one of the keychains in the **authentication-key-chains** statement at the **[edit security]** hierarchy level.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD Authentication for PIM on page 108
- *Understanding Bidirectional Forwarding Detection Authentication for PIM*
- authentication on page 181

local

Syntax

```
local {
  disable;
  address address;
  family (inet | inet6) {
    disable;
    address address;
    anycast-pim {
      local-address address;
        rp-set {
        address address <forward-msdp-sa>;
      }
    }
    group-ranges {
      destination-ip-prefix</prefix-length>;
    }
    hold-time seconds;
    override;
    priority number;
  }
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  hold-time seconds;
  override;
  priority number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
The remaining statements are explained separately.

Description

Configure the routing device's RP properties.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs on page 75

local-address (Protocols PIM)

Syntax local-address *address*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6) anycast-pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6) anycast-pim],
[edit protocols pim rp local family (inet | inet6) anycast-pim],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6) anycast-pim]

Release Information
Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the routing device local address for the anycast rendezvous point (RP). If this statement is omitted, the router ID is used as this address.

Options ***address***—Anycast RP IPv4 or IPv6 address, depending on **family** configuration.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring PIM Anycast With or Without MSDP on page 79

loose-check

Syntax loose-check;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection authentication],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection authentication]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD Authentication for PIM on page 108
- *Understanding Bidirectional Forwarding Detection Authentication for PIM*
- authentication on page 181

maximum-rps

Syntax maximum-rps *limit*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp embedded-rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp embedded-rp],
[edit protocols pim rp embedded-rp],
[edit routing-instances *routing-instance-name* protocols pim rp embedded-rp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Limit the number of RPs that the routing device acknowledges.

Options ***limit***—Number of RPs.

Range: 1 through 500

Default: 100

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring PIM Embedded RP for IPv6*

minimum-interval (PIM BFD Liveness Detection)

Syntax minimum-interval *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

Options ***milliseconds***—Minimum transmit and receive interval.

Range: 1 through 255,000 milliseconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 107

minimum-interval (PIM BFD Transmit Interval)

Syntax minimum-interval *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection transmit-interval],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection transmit-interval]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for BFD authentication introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the minimum interval after which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit protocols pim interface *interface-name* bfd-liveness-detection]** hierarchy level.

Options ***milliseconds***—Minimum transmit interval value.

Range: 1 through 255,000

NOTE: The threshold value specified in the **threshold** statement must be greater than the value specified in the **minimum-interval** statement for the **transmit-interval** statement.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 107
- bfd-liveness-detection on page 182
- minimum-interval on page 210
- threshold on page 236

minimum-receive-interval

Syntax minimum-receive-interval *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit protocols pim interface *interface-name* bfd-liveness-detection]** hierarchy level.

Options ***milliseconds***—Minimum receive interval.

Range: 1 through 255,000 milliseconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 107

mode (Protocols PIM)

Syntax mode (dense | sparse | sparse-dense);

Hierarchy Level [edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure PIM to operate in sparse, dense, or sparse-dense mode.

NOTE: The QFX Series does not support dense or sparse-dense mode.

Options **dense**—Operate in dense mode.

sparse—Operate in sparse mode.

sparse-dense—Operate in sparse-dense mode.

Default: **sparse**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

multiplier

Syntax multiplier *number*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Options *number*—Number of hello packets.

Range: 1 through 255

Default: 3

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 107

neighbor-policy

Syntax neighbor-policy [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply a PIM interface-level policy to filter neighbor IP addresses.

Options ***policy-name***—Name of the policy that filters neighbor IP addresses.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Interface-Level PIM Neighbor Policies on page 93

next-hop (PIM RPF Selection)

Syntax next-hop *next-hop-address*;

Hierarchy Level
[edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address* source *source-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address* wildcard-source],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses* source *source-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses* wildcard-source]

Release Information
Statement introduced in JUNOS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the specific next-hop address for the PIM group source.

Options ***next-hop-address***—Specific next-hop address for the PIM group source.

Required Privilege Level
view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation
- *Example: Configuring PIM RPF Selection*

no-adaptation (PIM BFD Liveness Detection)

Syntax no-adaptation;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in Junos OS Release 9.0
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for BFD authentication introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure BFD sessions not to adapt to changing network conditions. We recommend that you *do not* disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 107
- bfd-liveness-detection on page 182

override-interval

Syntax override-interval *milliseconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim]
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Set the maximum time in milliseconds to delay sending override join messages for a multicast network that has join suppression enabled. When a router or switch sees a prune message for a join it is currently suppressing, it waits for the interval specified by the override timer before it sends an override join message.

Options This is a random timer with a value in milliseconds.

Range: 0 through maximum override value

Default: 2000 milliseconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Enabling Join Suppression on page 69
- propagation-delay on page 224
- reset-tracking-bit on page 225

pim

Syntax

```
pim {
    disable;
    assert-timeout seconds;
    dense-groups {
        addresses;
    }
    dr-election-on-p2p;
    export;
    family (inet | inet6) {
        disable;
    }
    graceful-restart {
        disable;
        restart-duration seconds;
    }
    import [ policy-names ];
    interface interface-name {
        accept-remote-source;
        disable;
        family (inet | inet6) {
            disable;
        }
        hello-interval seconds;
        mode (dense | sparse | sparse-dense);
        neighbor-policy [ policy-names ];
        override-interval milliseconds;
        priority number;
        propagation-delay milliseconds;
        reset-tracking-bit;
        version version;
    }
    join-load-balance;
    join-prune-timeout;
    nonstop-routing;
    override-interval milliseconds;
    propagation-delay milliseconds;
    reset-tracking-bit;
    rib-group group-name;
    rp {
        auto-rp {
            (announce | discovery | mapping);
            (mapping-agent-election | no-mapping-agent-election);
        }
        bootstrap {
            family (inet | inet6) {
                export [ policy-names ];
                import [ policy-names ];
                priority number;
            }
        }
        bootstrap-import [ policy-names ];
        bootstrap-export [ policy-names ];
```

```
    bootstrap-priority number;
    dr-register-policy [ policy-names ];
    embedded-rp {
      group-ranges {
        destination-ip-prefix</prefix-length>;
      }
      maximum-rps limit;
    }
    local {
      family (inet | inet6) {
        address address;
        anycast-pim {
          disable;
            rp-set {
            address address <forward-msdp-sa>;
          }
          local-address address;
        }
        group-ranges {
          destination-ip-prefix</prefix-length>;
        }
        hold-time seconds;
        priority number;
      }
    }
    rp-register-policy [ policy-names ];
    spt-threshold {
      infinity [ policy-names ];
    }
    static {
      address address {
        group-ranges {
        version version;
        destination-ip-prefix</prefix-length>;
        }
      }
    }
  }
  rpf-selection {
    group group-address{
      source source-address{
        next-hop next-hop-address;
      }
      wildcard-source {
        next-hop next-hop-address;
      }
    }
    prefix-list prefix-list-addresses {
      source source-address {
        next-hop next-hop-address;
      }
      wildcard-source {
        next-hop next-hop-address;
      }
    }
  traceoptions {
```

```
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  tunnel-devices [ mt-fpc/pic/port ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information

Statement introduced before Junos OS Release 7.4.
family statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Enable PIM on the routing device.

The statements are explained separately.

Default

PIM is disabled on the routing device.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

prefix-list (PIM RPF Selection)

Syntax

```
prefix-list prefix-list-addresses {
  source source-address {
    next-hop next-hop-address;
  }
  wildcard-source {
    next-hop next-hop-address;
  }
}
```

Hierarchy Level

[edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address* source *source-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address* wildcard-source],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses* source *source-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses* wildcard-source]

Release Information

Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

(Optional) Configure a list of prefixes (addresses) for multiple PIM groups.

Options

prefix-list-addresses—List of prefixes (addresses) for multiple PIM groups.

The remaining statements are explained separately.

Required Privilege Level

view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- *Example: Configuring PIM RPF Selection*

priority (Bootstrap)

Syntax priority *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp bootstrap (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)],
[edit protocols pim rp bootstrap (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)]

Release Information
Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the routing device's likelihood to be elected as the bootstrap router.

Options
number—Routing device's priority for becoming the bootstrap router. A higher value corresponds to a higher priority.

Range: 0 through a 32-bit number

Default: 0 (The routing device has the least likelihood of becoming the bootstrap router and sends packets with a priority of 0.)

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Configuring PIM Bootstrap Properties for IPv4*
- *Configuring PIM Bootstrap Properties for IPv4 or IPv6*
- bootstrap-priority on page 186

priority (PIM Interfaces)

Syntax priority *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the routing device's likelihood to be elected as the designated router.

Options
number—Routing device's priority for becoming the designated router. A higher value corresponds to a higher priority.

Range: 0 through 4294967295

Default: 1 (Each routing device has an equal probability of becoming the DR.)

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Interface Priority for PIM Designated Router Selection on page 63

priority (PIM RPs)

Syntax priority *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp bidirectional address *address*],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit protocols pim rp bidirectional address *address*],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional RP addresses introduced in Junos OS Release 12.1.

Description For PIM-SM, configure this routing device's priority for becoming an RP.

For bidirectional PIM, configure this RP address' priority for becoming an RP.

The bootstrap router uses this field when selecting the list of candidate rendezvous points to send in the bootstrap message. A smaller number increases the likelihood that the routing device or RP address becomes the RP. A priority value of 0 means that bootstrap router can override the group range being advertised by the candidate RP.

Options ***number***—Priority for becoming an RP. A lower value corresponds to a higher priority.

Range: 0 through 255

Default: 1

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs on page 75 in the *Multicast Protocols Feature Guide for Routing Devices*
- *Example: Configuring Bidirectional PIM*

propagation-delay

Syntax propagation-delay *milliseconds*;

Hierarchy Level [edit protocols pim],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Set a delay for implementing a PIM prune message on the upstream routing device on a multicast network for which join suppression has been enabled. The routing device waits for the prune pending period to detect whether a join message is currently being suppressed by another routing device.

Options ***milliseconds***—Interval for the prune pending timer, which is the sum of the **propagation-delay** value and the **override-interval** value.

Range: 250 through 2000 milliseconds

Default: 500 milliseconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Enabling Join Suppression on page 69
- override-interval on page 216
- reset-tracking-bit on page 225

reset-tracking-bit

Syntax reset-tracking-bit;

Hierarchy Level [edit protocols pim],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Change the value of a tracking bit (T-bit) field in the LAN prune delay hello option from the default of 1 to 0, which enables join suppression for a multicast interface. When the network starts receiving multiple identical join messages, join suppression triggers a random timer with a value of 66 through 84 milliseconds (1.1 × periodic through 1.4 × periodic, where periodic is 60 seconds). This creates an interval during which no identical join messages are sent. Eventually, only one of the identical messages is sent. Join suppression is triggered each time identical messages are sent for the same join.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Enabling Join Suppression on page 69
- override-interval on page 216
- propagation-delay on page 224

rib-group (Protocols PIM)

Syntax

```
rib-group {
    inet group-name;
    inet6 group-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Associate a routing table group with PIM.

Options

table-name—Name of the routing table. The name must be one that you defined with the **rib-groups** statement at the **[edit routing-options]** hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring a Dedicated PIM RPF Routing Table*

rp

Syntax

```
rp {
    auto-rp {
        (announce | discovery | mapping);
        (mapping-agent-election | no-mapping-agent-election);
    }
    bidirectional {
        address address {
            group-ranges {
                destination-ip-prefix</prefix-length>;
            }
            hold-time seconds;
            priority number;
            }
        }
    bootstrap {
        family (inet | inet6) {
            export [ policy-names ];
            import [ policy-names ];
            priority number;
        }
    }
    bootstrap-export [ policy-names ];
    bootstrap-import [ policy-names ];
    bootstrap-priority number;
    dr-register-policy [ policy-names ];
    embedded-rp {
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
        maximum-rps limit;
    }
    group-rp-mapping {
        family (inet | inet6) {
            log-interval seconds;
            maximum limit;
            threshold value;
            }
        }
        log-interval seconds;
        maximum limit;
        threshold value;
        }
    }
    local {
        family (inet | inet6) {
            disable;
            address address;
            anycast-pim {
                local-address address;
                    address address <forward-msdp-sa>;
                    rp-set {
                }
```

```
        }
        group-ranges {
          destination-ip-prefix</prefix-length>;
        }
        hold-time seconds;
        override;
        priority number;
      }
    }
    register-limit {
      family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
        }
      }
      log-interval seconds;
      maximum limit;
      threshold value;
      }
    }
    register-probe-time register-probe-time;
    }
    rp-register-policy [ policy-names ];
    static {
      address address {
        override;
        version version;
        group-ranges {
          destination-ip-prefix</prefix-length>;
        }
      }
    }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure the routing device as an actual or potential RP. A routing device can be an RP for more than one group.

The remaining statements are explained separately.

Default

If you do not include the **rp** statement, the routing device can never become the RP.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Understanding PIM Sparse Mode on page 7

rp-register-policy

Syntax
rp-register-policy [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information
Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description
Apply one or more policies to control incoming PIM register messages.

Options
policy-names—Name of one or more import policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Register Message Filters on a PIM RP and DR on page 96
- dr-register-policy on page 189

rp-set

Syntax

```
rp-set {
    address address <forward-msdp-sa>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim local family (inet | inet6) anycast-pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim local family (inet | inet6) anycast-pim],
[edit protocols pim local family (inet | inet6) anycast-pim],
[edit routing-instances *routing-instance-name* protocols pim local family (inet | inet6) anycast-pim]

Release Information

Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PIM Anycast With or Without MSDP on page 79

rpf-selection

Syntax

```
rpf-selection {
  group group-address {
    sourcesource-address {
      next-hop next-hop-address;
    }
    wildcard-source {
      next-hop next-hop-address;
    }
  }
  prefix-list prefix-list-addresses {
    source source-address {
      next-hop next-hop-address;
    }
    wildcard-source {
      next-hop next-hop-address;
    }
  }
```

Hierarchy Level [edit routing-instances *routing-instance-name* protocols pim]

Release Information Statement introduced in JUNOS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the PIM RPF next-hop neighbor for a specific group and source for a VRF routing instance.

The remaining statements are explained separately.

Default If you omit the **rpf-selection** statement, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows.

Options ***source-address***—Specific source address for the PIM group.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- *Example: Configuring PIM RPF Selection*

source (PIM RPF Selection)

Syntax

```
source source-address {
    next-hop next-hop-address;
}
```

Hierarchy Level

[edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses*]

Release Information

Statement introduced in JUNOS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure the source address for the PIM group.

Options

source-address—Specific source address for the PIM group.

The remaining statements are explained separately.

Required Privilege Level

view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- *Example: Configuring PIM RPF Selection*

spt-threshold

Syntax

```
spt-threshold {
  infinity [ policy-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Set the SPT threshold to infinity for a source-group address pair. Last-hop multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or an SPT rooted at the source. By default, last-hop routing devices transition to a direct SPT to the source. You can configure this routing device to set the SPT transition value to infinity to prevent this transition for any source-group address pair.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring the PIM SPT Threshold Policy on page 101

static (Protocols PIM)

Syntax

```
static {
  address address {
    group-ranges {
      destination-ip-prefix</prefix-length>;
    }
    override;
    version version;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure static RP addresses. The default static RP address is 224.0.0.0/4. To configure other addresses, include one or more **address** statements. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.

For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Static PIM RP Address on the Non-RP Routing Device on page 77

threshold (PIM BFD Detection Time)

Syntax threshold *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection detection-time],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection detection-time]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for BFD authentication introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

NOTE: The threshold value must be equal to or greater than the transmit interval.

The threshold time must be equal to or greater than the value specified in the minimum-interval or the minimum-receive-interval statement.

Options *milliseconds*—Value for the detection time adaptation threshold.

Range: 1 through 255,000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 107
- bfd-liveness-detection on page 182
- detection-time on page 187
- minimum-interval on page 210
- minimum-receive-interval on page 212

threshold (PIM BFD Transmit Interval)

Syntax threshold *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection transmit-interval],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection transmit-interval]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.

Options ***milliseconds***—Value for the transmit interval adaptation threshold.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

NOTE: The threshold value specified in the threshold statement must be greater than the value specified in the minimum-interval statement for the transmit-interval statement.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 107
- bfd-liveness-detection on page 182

transmit-interval (PIM BFD Liveness Detection)

Syntax

```
transmit-interval {
  minimum-interval milliseconds;
  threshold milliseconds;
}
```

Hierarchy Level

[edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for BFD authentication introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Specify the transmit interval for the **bfd-liveness-detection** statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 107
- bfd-liveness-detection on page 182
- threshold on page 236
- minimum-interval on page 211
- minimum-receive-interval on page 212

traceoptions (Protocols PIM)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure PIM tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

Default

The default PIM trace options are those inherited from the routing protocol's **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the **pim-log** file.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

Range: 2 through 1000 files

Default: 2 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

PIM Tracing Flags

- **assert**—Assert messages
- **bidirectional-df-election**—Bidirectional PIM designated-forwarder (DF) election events

- **bootstrap**—Bootstrap messages
- **cache**—Packets in the PIM sparse mode routing cache
- **graft**—Graft and graft acknowledgment messages
- **hello**—Hello packets
- **join**—Join messages
- **mt**—Multicast tunnel messages
- **nsr-synchronization**—Nonstop active routing (NSR) synchronization messages
- **packets**—All PIM packets
- **prune**—Prune messages
- **register**—Register and register stop messages
- **rp**—Candidate RP advertisements
- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 0 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Trace Options on page 57
- *Tracing DVMRP Protocol Traffic*
- Tracing MSDP Protocol Traffic on page 147
- Configuring PIM Trace Options on page 57

version (BFD)

Syntax version (0 | 1 | automatic);

Hierarchy Level [edit protocols piminterface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the bidirectional forwarding detection (BFD) protocol version that you want to detect.

Options Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version)

Default: automatic

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 107

version (PIM)

Syntax version *version*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* protocols pim rp static address *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp static address *address*],
[edit protocols pim interface *interface-name*],
[edit protocols pim rp static address *address*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim rp static address *address*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Specify the version of PIM.

Options ***version***—PIM version number.

Range: 1 or 2

Default: PIMv1 for rendezvous point (RP) mode (at the **[edit protocols pim rp static address *address*]** hierarchy level). PIMv2 for interface mode (at the **[edit protocols pim interface *interface-name*]** hierarchy level).

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling PIM Sparse Mode on page 65
- *Configuring PIM Dense Mode Properties*
- *Configuring PIM Sparse-Dense Mode Properties*

wildcard-source (PIM RPF Selection)

Syntax

```
wildcard-source {
    next-hop next-hop-address;
}
```

Hierarchy Level

[edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses*]

Release Information

Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Use a wildcard for the multicast source instead of (or in addition to) a specific multicast source.

The remaining statements are explained separately.

Required Privilege Level

view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- *Example: Configuring PIM RPF Selection*

CHAPTER 28

IGMP Configuration Statements

- traceoptions (Protocols IGMP) on page 267
- version (Protocols IGMP) on page 269

accounting (Protocols IGMP)

Syntax accounting;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Enable the collection of IGMP join and leave event statistics on the system.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Recording IGMP Join and Leave Events on page 130

accounting (Protocols IGMP Interface)

Syntax (accounting | no-accounting);

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Enable or disable the collection of IGMP join and leave event statistics for an interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Recording IGMP Join and Leave Events on page 130

asm-override-ssm

Syntax asm-override-ssm;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Enable the routing device to accept any-source multicast join messages (*,G) for group addresses that are within the default or configured range of source-specific multicast groups.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 165

disable (Protocols IGMP)

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Disable IGMP on the system.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling IGMP on page 134

exclude (Protocols IGMP)

Syntax exclude;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the static group to operate in exclude mode. In exclude mode all sources except the address configured are accepted for the group. If this statement is not included, the group operates in include mode.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 123

group (Protocols IGMP)

Syntax

```
group multicast-group-address {
  exclude;
  group-count number;
  group-increment increment;
  source ip-address {
    source-count number;
    source-increment increment;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static], [edit protocols igmp interface *interface-name* static]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.

NOTE: You must specify a unique address for each group.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 123

group-count (Protocols IGMP)

Syntax group-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the number of static groups to be created.

Options *number*—Number of static groups.

Default:

Range: 1 through 512

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 123

group-increment (Protocols IGMP)

Syntax group-increment *increment*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the number of times the address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.

Options *increment*—Number of times the address should be incremented.

Default: 0.0.0.1

Range: 0.0.0.1 through 255.255.255.255

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 123

group-limit

Syntax group-limit *limit*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*], [edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.

To confirm the configured group limit on the interface, use the show igmp interface command.

Default By default, there is no limit to the number of multicast groups that can join the interface.

Options ***limit***—group limit value for the interface.

Range: 1 through 32767

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 131
- *group-threshold*
- *log-interval*

group-policy (Protocols IGMP)

Syntax group-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 119

igmp

Syntax

```
igmp {
  accounting;
  interface interface-name {
    disable;
    (accounting | no-accounting);
    group-limit limit;
    group-policy [ policy-names ];
    group-threshold
    immediate-leave;
    log-interval
    oif-map map-name;
    passive;
    promiscuous-mode;
    ssm-map ssm-map-name;
    ssm-map-policy ssm-map-policy-name;
    static {
      group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
          source-count number;
          source-increment increment;
        }
      }
    }
    version version;
  }
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Enable IGMP on the router or switch. IGMP must be enabled for the router or switch to receive multicast packets.

The remaining statements are explained separately.

Default IGMP is disabled on the router or switch. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling IGMP on page 115

immediate-leave (Protocols IGMP)

Syntax immediate-leave;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*], [edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

Starting in Junos OS Release 9.3, both IGMP version 2 and IGMP version 3 do host tracking when the **immediate-leave** statement is configured. This means that the multicast group leaves only when the last host leaves. The routing device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.

NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the routing device only knows about the one interested host and does not know about the others.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Specifying Immediate-Leave Host Removal for IGMP on page 118

interface (Protocols IGMP)

Syntax

```
interface interface-name {
  disable;
  (accounting | no-accounting);
  group-limit limit;
  group-policy [ policy-names ];
  immediate-leave;
  oif-map map-name;
  passive;
  promiscuous-mode;
  ssm-map ssm-map-name;
  ssm-map-policy ssm-map-policy-name;
  static {
    group multicast-group-address {
      exclude;
      group-count number;
      group-increment increment;
      source ip-address {
        source-count number;
        source-increment increment;
      }
    }
  }
  version version;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Enable IGMP on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling IGMP on page 115

maximum-transmit-rate (Protocols IGMP)

Syntax maximum-transmit-rate *packets-per-second*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Limit the transmission rate of IGMP packets

Options **packets-per-second**—Maximum number of IGMP packets transmitted in one second by the routing device.

Range: 1 through 10000

Default: 500 packets

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Limiting the Maximum IGMP Message Rate on page 123

oif-map

Syntax oif-map *map-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Associates an outgoing interface (OIF) map to the IGMP interface. The OIF map is a routing policy statement that can contain multiple terms.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Multicast with Subscriber VLANs*

passive (IGMP)

Syntax passive <allow-receive> <send-general-query> <send-group-query>;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*], [edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6.
allow-receive, **send-general-query**, and **send-group-query** options were added in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify that IGMP run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as IGMP reports, queries, and leaves.

NOTE: You can selectively activate up to two out of the three available options for the passive statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the passive statement.

Options

allow-receive—Enables IGMP to receive control traffic on the interface.

send-general-query—Enables IGMP to send general queries on the interface.

send-group-query—Enables IGMP to send group-specific and group-source-specific queries on the interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Multicast with Subscriber VLANs*
- Enabling IGMP on page 115

promiscuous-mode (Protocols IGMP)

Syntax promiscuous-mode;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 9.2 for dynamic profiles.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify that the interface accepts IGMP reports from hosts on any subnetwork. Note that when enabling promiscuous-mode, all routing devices on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Configuring a Dynamic Profile for Client Access*
- Accepting IGMP Messages from Remote Subnetworks on page 120

query-interval (Protocols IGMP)

Syntax query-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify how often the querier routing device sends general host-query messages.

Options *seconds*—Time interval.
Range: 1 through 1024
Default: 125 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Modifying the IGMP Host-Query Message Interval on page 117
- query-last-member-interval (Protocols IGMP) on page 260
- query-response-interval (Protocols IGMP) on page 261

query-last-member-interval (Protocols IGMP)

Syntax query-last-member-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify how often the querier routing device sends group-specific query messages.

Options ***seconds***—Time interval, in fractions of a second or seconds.

Range: 0.1 through 0.9, then in 1-second intervals 1 through 999999

Default: 1 second

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the IGMP Last-Member Query Interval on page 117
- query-interval (Protocols IGMP) on page 259
- query-response-interval (Protocols IGMP) on page 261

query-response-interval (Protocols IGMP)

Syntax query-response-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify how long the querier routing device waits to receive a response to a host-query message from a host.

Options ***seconds***—The query response interval must be less than the query interval.

Range: 1 through 1024

Default: 10 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the IGMP Query Response Interval on page 121
- query-interval (Protocols IGMP) on page 259
- query-last-member-interval (Protocols IGMP) on page 260

robust-count (Protocols IGMP)

Syntax robust-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Tune the expected packet loss on a subnet. This factor is used to calculate the group member interval, other querier present interval, and last-member query count.

Options *number*—Robustness variable.

Range: 2 through 10

Default: 2

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the IGMP Robustness Variable on page 122

source (Protocols IGMP)

Syntax

```
source ip-address {
    source-count number;
    source-increment increment;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Specify the IP version 4 (IPv4) unicast source address for the multicast group being statically configured on an interface.

Options

ip-address—IPv4 unicast address.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 123

source-count (Protocols IGMP)

Syntax source-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address* source],
[edit protocols igmp interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the number of multicast source addresses that should be accepted for each static group created.

Options *number*—Number of source addresses.

Default: 1

Range: 1 through 1024

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 123

source-increment (Protocols IGMP)

Syntax source-increment *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address* source],
[edit protocols igmp interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the number of times the multicast source address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.

Options *increment*—Number of times the source address should be incremented.

Default: 0.0.0.1

Range: 0.0.0.1 through 255.255.255.255

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 123

ssm-map (Protocols IGMP)

Syntax ssm-map *ssm-map-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Apply an SSM map to an IGMP interface.

Options ***ssm-map-name***—Name of SSM map.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring SSM Mapping on page 163

ssm-map-policy (IGMP)

Syntax ssm-map-policy *ssm-map-policy-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Apply an SSM map policy to an IGMP interface.

Options ***ssm-map-policy-name***—Name of SSM map policy.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring SSM Maps for Different Groups to Different Sources on page 169

static (Protocols IGMP)

Syntax

```
static {
  group multicast-group-address {
    exclude;
    group-count number;
    group-increment increment;
    source ip-address {
      source-count number;
      source-increment increment;
    }
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*], [edit protocols igmp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Test multicast forwarding on an interface without a receiver host.

The **static** statement simulates IGMP joins on a routing device statically on an interface without any IGMP hosts. It is supported for both IGMPv2 and IGMPv3 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.

NOTE: To prevent joining too many groups accidentally, the **static** statement is not supported with the **interface all** statement.

The remaining statements are explained separately.

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 123

traceoptions (Protocols IGMP)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure IGMP tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

To trace the paths of multicast packets, use the **mtrace** command.

Default

The default IGMP trace options are those inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the file **igmp-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

Range: 2 through 1000 files

Default: 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

IGMP Tracing Flags

- **leave**—Leave group messages (for IGMP version 2 only).
- **mtrace**—Mtrace packets. Use the **mtrace** command to troubleshoot the software.
- **packets**—All IGMP packets.

- **query**—IGMP membership query messages, including general and group-specific queries.
- **report**—Membership report messages.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level
routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation
- Tracing IGMP Protocol Traffic on page 132

version (Protocols IGMP)

Syntax
version *version*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description
Specify the version of IGMP.

Options
version—IGMP version number.

Range: 1, 2, or 3

Default: IGMP version 2

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Changing the IGMP Version on page 116

CHAPTER 29

IGMP Snooping Configuration Statements

data-forwarding

Syntax

```
data-forwarding {
  receiver {
    source-vlans vlan-list;
    install;
  }
  source {
    groups group-prefix;
  }
}
```

Hierarchy Level [edit protocols igmp-snooping vlan (all | *vlan-name*)]

Release Information Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for the QFX Series.

Description Configure the VLAN to be a multicast source VLAN (MVLAN) or a multicast VLAN registration (MVR) receiver VLAN. Each data-forwarding VLAN, which can be a multicast source VLAN (MVLAN) or a multicast receiver VLAN, must have exactly one source statement or exactly one receiver statement. A data-forwarding VLAN can operate only in IGMP version 2 (IGMPv2) mode.

The remaining statements are explained separately.

Default Disabled

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Multicast VLAN Registration on page 140
- Configuring Multicast VLAN Registration (CLI Procedure) on page 139

disable (IGMP Snooping)

Syntax

```
disable;
```

Hierarchy Level [edit protocols igmp-snooping vlan *vlan-name*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Disable IGMP snooping on all interfaces in a VLAN.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135

group (IGMP Snooping)

Syntax group *ip-address*;

Hierarchy Level [edit protocols igmp-snooping vlan *vlan-name* interface *interface-name* static]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure a static multicast group using a valid IP multicast address.

Default None.

Options ***ip-address*** —IP address of the multicast group receiving data on an interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- show igmp-snooping vlans on page 371
- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135

groups (Multicast VLAN Registration)

Syntax groups *group-prefix*;

Hierarchy Level [edit protocols igmp-snooping vlan (all | *vlan-name*) data-forwarding source]

Release Information Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for the QFX Series.

Description Specify the IP address range of the multicast VLAN (MVLAN) source interfaces.

Default Disabled

Options *group-prefix*—IP address range of the source group. Each MVLAN must have exactly one **groups** statement. If there are multiple MVLANs on the switch, their group ranges must be unique.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Multicast VLAN Registration on page 140
- Configuring Multicast VLAN Registration (CLI Procedure) on page 139

igmp-snooping

Syntax

```
igmp-snooping {
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
      regex>;
    flag flag (detail | disable | receive | send);
  }
  vlan vlan-name {
    data-forwarding {
      source {
        groups group-prefix;
      }
      receiver {
        source-vlans vlan-list;
        install ;
      }
    }
    disable;
    immediate-leave;
    interface interface-name {
      multicast-router-interface;
      static {
        group ip-address;
      }
    }
    robust-count number;
    version number;
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
version statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Enable and configure IGMP snooping.

The remaining statements are explained separately.

Default IGMP snooping is disabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135

install (Multicast VLAN Registration)

Syntax install;

Hierarchy Level [edit protocols igmp-snooping vlan (all | *vlan-name*) data-forwarding receiver]

Release Information Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for the QFX Series.

Description Install forwarding entries in the multicast receiver VLAN. By default, the multicast VLAN (MVLAN) installs forwarding entries for MVLAN groups only.

Default Disabled

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Multicast VLAN Registration on page 140
- Configuring Multicast VLAN Registration (CLI Procedure) on page 139

interface (IGMP Snooping)

Syntax
```
interface interface-name {
  multicast-router-interface;
  static {
    group ip-address;
  }
}
```

Hierarchy Level [edit protocols igmp-snooping vlan *vlan-name*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Enable IGMP snooping on an interface and configure interface-specific properties.

The remaining statements are explained separately.

Options *interface-name*—Name of the interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135
- show igmp-snooping vlans on page 371

multicast-router-interface (IGMP Snooping)

Syntax multicast-router-interface;

Hierarchy Level [edit protocols igmp-snooping vlan *vlan-name* interface *interface-name*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure an interface to forward IGMP messages to multicast routers.

Default Disabled. If this statement is disabled, the interface drops IGMP messages it receives.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- show igmp-snooping vlans on page 371
- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135

proxy (Multicast VLAN Registration)

Syntax proxy source-address *ip-address*;

Hierarchy Level [edit protocols igmp-snooping vlan (all | *vlan-name*)]

Release Information Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for the QFX Series.

Description Specify that the VLAN operate in proxy mode. The proxy option is supported only for a VLAN acting as a data-forwarding source.

Default Disabled

Options **source-address *ip-address***—IP address of the source VLAN to act as proxy.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Multicast VLAN Registration on page 140
- Configuring Multicast VLAN Registration (CLI Procedure) on page 139

receiver

Syntax

```
receiver {
  source-vlans vlan-list;
  install;
}
```

Hierarchy Level [edit protocols igmp-snooping vlan (all | *vlan-name*) data-forwarding]

Release Information Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for the QFX Series.

Description Configure a VLAN as a multicast receiver VLAN of the multicast VLAN (MVLAN).

The remaining statements are explained separately.

Default Disabled

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Multicast VLAN Registration on page 140
- Configuring Multicast VLAN Registration (CLI Procedure) on page 139

robust-count (IGMP Snooping)

Syntax robust-count *number*;

Hierarchy Level [edit protocols igmp-snooping vlan *vlan-name*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the number of intervals the switch waits before removing a multicast group from the multicast forwarding table. Configure the length of each interval using the **query-interval** statement.

Default 2 intervals

Options *number*—Number of intervals the switch waits before timing out a multicast group.

Range: 2 through 10

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135
- show igmp-snooping vlans on page 371

source (Multicast VLAN Registration)

Syntax

```
source {
  groups group-prefix;
}
```

Hierarchy Level [edit protocols igmp-snooping vlan (all | *vlan-name*) data-forwarding]

Release Information Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for the QFX Series.

Description Configure a VLAN to be a multicast source VLAN (MVLAN).

The remaining statement is explained separately.

Default Disabled

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Multicast VLAN Registration on page 140
- Configuring Multicast VLAN Registration (CLI Procedure) on page 139

source-vlans

Syntax source-vlans *vlan-list*;

Hierarchy Level [edit protocols igmp-snooping vlan (all | *vlan-name*) data-forwarding receiver]

Release Information Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for the QFX Series.

Description Specify a list of multicast VLANs (MVLANs) from which this multicast receiver VLAN receives multicast traffic. Either all of these MVLANs must be in proxy mode or none of them can be in proxy mode.

Default Disabled

Options *vlan-list*—Names of the MVLANs.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Multicast VLAN Registration on page 140
- Configuring Multicast VLAN Registration (CLI Procedure) on page 139

static (IGMP Snooping)

Syntax

```
static {
  group ip-address;
}
```

Hierarchy Level [edit protocols igmp-snooping vlan *vlan-name* interface *interface-name*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Statically define multicast groups on an interface.

The remaining statement is explained separately.

Default No multicast groups are statically defined.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135
- show igmp-snooping vlans on page 371

traceoptions (IGMP Snooping)

Syntax

```
traceoptions {
  file filename <files number> <no-stamp> <size size> <replace> <world-readable |
    no-world-readable>;
  flag flag (detail | disable | receive | send);
}
```

Hierarchy Level

For platforms without ELS:

[edit protocols igmp-snooping]

For platforms with ELS:

[edit protocols igmp-snooping vlan]

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Define tracing operations for IGMP snooping.

Default

The **traceoptions** feature is disabled by default.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached (***x***k to specify KB, ***x***m to specify MB, or ***x***g to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag* —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.
- **general**—Trace general IGMP snooping protocol events.
- **krt**—Trace communication over routing sockets.
- **nexthop**— Trace next-hop related events.
- **normal**—Trace normal IGMP snooping protocol events.
- **packets**—Trace all IGMP packets.
- **policy**—Trace policy processing.
- **query**—Trace IGMP membership query messages.
- **report**—Trace membership report messages.

- **route**—Trace routing information.
- **state**—Trace IGMP state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.
- **vlan**—Trace VLAN related events.

no-stamp—(Optional) Do not time stamp trace file.

no-world-readable—(Optional) Restrict file access to the user who created the file.

size *size* —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option. Use ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes.

Range: 10 KB through 1 gigabytes

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135

vlan (IGMP Snooping)

Syntax

```
vlan vlan-name {
  immediate-leave;
  interface interface-name {
    multicast-router-interface;
    static {
      group ip-address;
    }
  }
  version number;
}
```

Hierarchy Level [edit protocols igmp-snooping]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
version statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure IGMP snooping parameters for a VLAN.

The remaining statements are explained separately.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range. For IGMP snooping, secondary private VLANs are not listed.

Default IGMP snooping options apply to the specified VLAN.

Options ***vlan-name***—Name of a VLAN.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135
- show igmp-snooping vlans on page 371

version (IGMP Snooping)

Syntax version *number*;

Hierarchy Level [edit protocols igmp-snooping vlan (all | *vlan-name*)]

Release Information Statement introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the IGMP version for the IGMP general query that the switch sends to hosts when an interface comes up. The configured IGMP version affects only the version of the general queries sent by a switch. It does not affect the version of IGMP messages that the switch can snoop. For example, If the switch is configured for IGMP version 1 (IGMPv1), it can snoop IGMPv2 and IGMPv3 messages.

Default If you do not configure the **version** statement, the default is IGMPv2.

Options ***version***—IGMP version number.
Range: 1 and 2.

NOTE: **IGMP v3 snooping is not supported.**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring IGMP Snooping (CLI Procedure)*
- Configuring IGMP Snooping on page 135

CHAPTER 30

MSDP Configuration Statements

active-source-limit

Syntax

```
active-source-limit {
    log-interval seconds;
    log-warning value;
    maximum number;
    threshold number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* protocols msdp source *ip-address/prefix-length*],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp source *ip-address/prefix-length*],
[edit protocols msdp],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit protocols msdp source *ip-address/prefix-length*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp source *ip-address/prefix-length*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Limit the number of active source messages the routing device accepts.

Default

If you do not include this statement, the router accepts any number of MSDP active source messages.

Options

The options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 150

authentication-key

Syntax authentication-key *peer-key*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Associate a Message Digest 5 (MD5) signature option authentication key with an MSDP peering session.

Default If you do not include this statement, the routing device accepts any valid MSDP messages from the peer address.

Options ***peer-key***—MD5 authentication key. The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of (,), &, and [. If you include spaces in an MSDP authentication key, enclose all characters in quotation marks (“ ”).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring MSDP in a Routing Instance*

data-encapsulation

Syntax data-encapsulation (disable | enable);

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure a rendezvous point (RP) using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages.

Default If you do not include this statement, the RP encapsulates multicast data.

Options **disable**—(Optional) Do not use MSDP data encapsulation.

enable—Use MSDP data encapsulation.

Default: **enable**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 150

default-peer

Syntax default-peer;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Establish this peer as the default MSDP peer and accept source-active messages from the peer without the usual peer-reverse-path-forwarding (peer-RPF) check.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 150

disable (Protocols MSDP)

Syntax disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Explicitly disable MSDP.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Disabling MSDP*

export (Protocols MSDP)

Syntax export [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Apply one or more policies to routes being exported from the routing table into MSDP.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring MSDP in a Routing Instance*
- import on page 293

group

Syntax

```
group group-name {
  disable;
  export [ policy-names ];
  import [ policy-names ];
  local-address address;
  mode (mesh-group | standard);
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  peer address; {
    disable;
    active-source-limit {
      maximum number;
      threshold number;
    }
    authentication-key peer-key;
    default-peer;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Define an MSDP peer group. MSDP peers within groups share common tracing options, if present and not overridden for an individual peer with the peer statement. To configure multiple MSDP groups, include multiple **group** statements.

By default, the group's options are identical to the global MSDP options. To override the global options, include group-specific options within the **group** statement.

The group must contain at least one peer.

Options

group-name—Name of the MSDP group.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring MSDP in a Routing Instance*

import (Protocols MSDP)

Syntax
import [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description
Apply one or more policies to routes being imported into the routing table from MSDP.

Options
policy-names—Name of one or more policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring MSDP in a Routing Instance*
- export on page 291

local-address

Syntax local-address *address*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the local end of an MSDP session. You must configure at least one peer for MSDP to function. When configuring a peer, you must include this statement. This address is used to accept incoming connections to the peer and to establish connections to the remote peer.

Options ***address***—IP address of the local end of the connection.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring MSDP in a Routing Instance*

maximum

Syntax maximum *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp active-source-limit],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp active-source-limit],
[edit protocols msdp active-source-limit],
[edit routing-instances *routing-instance-name* protocols msdp active-source-limit]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the maximum number of MSDP active source messages the router accepts.

Options ***number***—Maximum number of active source messages.

Range: 1 through 1,000,000

Default: 25,000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 150
- threshold on page 302

mode (Protocols MSDP)

Syntax mode (mesh-group | standard);

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure groups of peers in a full mesh topology to limit excessive flooding of source-active messages to neighboring peers. The default flooding mode is **standard**.

Default If you do not include this statement, default flooding is applied.

Options **mesh-group**—Group of peers that are mesh group members.

standard—Use standard MSDP source-active flooding rules.

Default: standard

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 150

msdp

Syntax

```
msdp {
    disable;
    active-source-limit {
        log-interval seconds;
        log-warning value;
        maximum number;
        threshold number;
    }
    data-encapsulation (disable | enable);
    export [ policy-names ];
    group group-name {
        ... group-configuration ...
    }
    hold-time seconds;
    import [ policy-names ];
    local-address address;
    keep-alive seconds;
    peer address {
        ... peer-configuration ...
    }
    rib-group group-name;
    source ip-prefix</prefix-length> {
        active-source-limit {
            maximum number;
            threshold number;
        }
    }
    sa-hold-time seconds;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    group group-name {
        disable;
        export [ policy-names ];
        import [ policy-names ];
        local-address address;
        mode (mesh-group | standard);
        peer address {
            ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
                just following ...
        }
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
    }
    peer address {
        disable;
        active-source-limit {
            maximum number;
            threshold number;
```

```
    }
    authentication-key peer-key;
    default-peer;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.4 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Enable MSDP on the router or switch. You must also configure at least one peer for MSDP to function.

Default

MSDP is disabled on the router or switch.

Options

The statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring MSDP in a Routing Instance*

peer (Protocols MSDP)

Syntax

```
peer address {
  disable;
  active-source-limit {
    maximum number;
    threshold number;
  }
  authentication-key peer-key;
  default-peer;
  export [ policy-names ];
  import [ policy-names ];
  local-address address;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Define an MSDP peering relationship. An MSDP routing device must know which routing devices are its peers. You define the peer relationships explicitly by configuring the neighboring routing devices that are the MSDP peers of the local routing device. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. To configure multiple MSDP peers, include multiple **peer** statements.

By default, the peer's options are identical to the global or group-level MSDP options. To override the global or group-level options, include peer-specific options within the peer (Protocols MSDP) statement.

At least one peer must be configured for MSDP to function. You must configure ***address*** and **local-address**.

Options

address—Name of the MSDP peer.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring MSDP in a Routing Instance*

rib-group (Protocols MSDP)

Syntax rib-group *group-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Associate a routing table group with MSDP.

Options *group-name*—Name of the routing table group. The name must be one that you defined with the **rib-groups** statement at the **[edit routing-options]** hierarchy level.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring MSDP in a Routing Instance*

source

Syntax

```
source ip-address</prefix-length> {
  active-source-limit {
    maximum number;
    threshold number;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Limit the number of active source messages the routing device accepts from sources in this address range.

Default

If you do not include this statement, the routing device accepts any number of MSDP active source messages.

Options

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 150

threshold

Syntax threshold *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp active-source-limit],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp active-source-limit],
[edit protocols msdp active-source-limit],
[edit routing-instances *routing-instance-name* protocols msdp active-source-limit]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the random early detection (RED) threshold for MSDP active source messages. This number must be less than the configured or default maximum.

Options *number*—RED threshold for active source messages.

Range: 1 through 1,000,000

Default: 24,000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 150
- maximum on page 295

traceoptions (Protocols MSDP)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure MSDP tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

Default

The default MSDP trace options are those inherited from the routing protocol's **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the **msdp-log** file.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

Range: 2 through 1000 files

Default: 2 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

MSDP Tracing Flags

- **keepalive**—Keepalive messages
- **packets**—All MSDP packets
- **route**—MSDP changes to the routing table
- **source-active**—Source-active packets
- **source-active-request**—Source-active request packets
- **source-active-response**—Source-active response packets

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow any user to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*****.0**. When ***trace-file*** again reaches this size, ***trace-file*****.0** is renamed ***trace-file*****.1** and ***trace-file*** is renamed ***trace-file*****.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: ***x*****k** to specify KB, ***x*****m** to specify MB, or ***x*****g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Tracing MSDP Protocol Traffic on page 147

CHAPTER 31

Source-Specific Multicast Configuration Statements

asm-override-ssm

Syntax asm-override-ssm;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Enable the routing device to accept any-source multicast join messages (*,G) for group addresses that are within the default or configured range of source-specific multicast groups.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 165

policy (SSM Maps)

Syntax policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-options multicast ssm-map *ssm-map-name*]

Release Information Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Apply one or more policies to an SSM map.

Options ***policy-names***—Name of one or more policies for SSM mapping.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To view this statement in the configuration.

Related Documentation
- Example: Configuring SSM Mapping on page 163

ssm-groups

Syntax ssm-groups [*ip-addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure source-specific multicast (SSM) groups.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the **ssm-groups** statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the **ssm-groups** statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.

IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.

Options ***ip-addresses***—List of one or more additional SSM group addresses separated by a space.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 165

ssm-map (Protocols IGMP)

Syntax ssm-map *ssm-map-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Apply an SSM map to an IGMP interface.

Options ***ssm-map-name***—Name of SSM map.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring SSM Mapping on page 163

ssm-map (Routing Options Multicast)

Syntax
```
ssm-map ssm-map-name {
  policy [ policy-names ];
  source [ addresses ];
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure SSM mapping.

Options ***ssm-map-name***—Name of the SSM map.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring SSM Mapping on page 163

ssm-map-policy (IGMP)

Syntax ssm-map-policy *ssm-map-policy-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Apply an SSM map policy to an IGMP interface.

Options ***ssm-map-policy-name***—Name of SSM map policy.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring SSM Maps for Different Groups to Different Sources on page 169

PART 3

Administration

CHAPTER 32

Routine Monitoring

Monitoring IGMP Snooping

Purpose Use the monitoring feature to view status and information about the IGMP snooping configuration.

Action To display IGMP snooping details in the CLI, enter the following commands:

- **show igmp-snooping vlans**
- **show igmp-snooping statistics**
- **show igmp-snooping route**
- **show igmp-snooping membership**

Meaning Table 8 on page 315 summarizes the IGMP snooping details displayed.

Table 8: Summary of IGMP Snooping Output Fields

Field	Values
IGMP Snooping Monitor	
VLAN	VLAN for which IGMP snooping is enabled.
Interfaces	Interface connected to a multicast router.
Groups	Number of the multicast groups learned by the VLAN.
MRouters	Multicast router.
Receivers	Multicast receiver.
IGMP Route Information	
VLAN	VLAN for which IGMP snooping is enabled.

Table 8: Summary of IGMP Snooping Output Fields *(continued)*

Field	Values
Next-Hop	Next hop assigned by the switch after performing the route lookup.
Group	Multicast groups learned by the VLAN.

Related Documentation

- IGMP Snooping Overview on page 35
- Example: Configuring IGMP Snooping on page 137
- Configuring IGMP Snooping on page 135
- Changing the IGMP Snooping Group Timeout Value on page 136

Verifying the IGMP Snooping Group Timeout Value

Purpose Verify that the IGMP snooping group timeout value has been changed correctly from its default value.

Action Display the IGMP snooping membership information, which contains the group timeout value that was derived from the IGMP configuration:

```
user@switch> show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
      Group: 225.0.0.1
      Receiver count: 1, Flags: <v2-hosts>
          ge-0/0/15.0 Uptime: 00:00:05 timeout: 510
```

Meaning The IGMP snooping group timeout value determines how long a switch waits to receive an IGMP query from a multicast router before removing a multicast group from its multicast cache table. When you enable IGMP snooping, the default IGMP snooping group timeout value of 260 seconds is applied to all VLANs, which means that the switch waits 260 seconds to receive an IGMP query before removing a multicast group from its multicast cache table. You can change the timeout value by using the **robust-count** option.

Related Documentation

- Changing the IGMP Snooping Group Timeout Value on page 136

CHAPTER 33

Monitoring Commands for Multicast Protocols

- clear igmp membership
- clear igmp-snooping membership
- clear igmp statistics
- clear igmp-snooping statistics
- clear msdp cache
- clear msdp statistics
- clear multicast bandwidth-admission
- clear multicast scope
- clear multicast sessions
- clear multicast statistics
- clear pim join
- clear pim register
- clear pim statistics
- mtrace
- mtrace from-source
- mtrace monitor
- mtrace to-gateway
- show configuration protocols igmp
- show igmp group
- show igmp interface
- show igmp statistics
- show igmp-snooping membership
- show igmp-snooping route
- show igmp-snooping statistics
- show igmp-snooping vlans
- show msdp

- show msdp source
- show msdp source-active
- show msdp statistics
- show multicast flow-map
- show multicast interface
- show multicast mrinfo
- show multicast next-hops
- show multicast pim-to-igmp-proxy
- show multicast pim-to-mld-proxy
- show multicast route
- show multicast rpf
- show multicast scope
- show multicast sessions
- show multicast usage
- show pim bootstrap
- show pim interfaces
- show pim join
- show pim neighbors
- show pim rps
- show pim source
- show pim statistics
- show system statistics igmp
- test msdp

clear igmp membership

List of Syntax
Syntax on page 319
Syntax (EX Series Switch and the QFX Series) on page 319

Syntax
clear igmp membership
<group *address-range*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
clear igmp membership
<group *address-range*>
<interface *interface-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Clear Internet Group Management Protocol (IGMP) group members.

Options
none—Clear all IGMP members on all interfaces and for all address ranges.

group *address-range*—(Optional) Clear all IGMP members that are in a particular address range. An example of a range is **224.2/16**. If you omit the destination prefix length, the default is **/32**.

interface *interface-name*—(Optional) Clear all IGMP group members on an interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- show igmp group on page 353
- show igmp interface on page 357

List of Sample Output
clear igmp membership on page 319
clear igmp membership interface on page 320
clear igmp membership group on page 321

Output Fields
See show igmp group for an explanation of output fields.

Sample Output

clear igmp membership

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

```
user@host> show igmp group
```

```
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          186
so-0/0/0         224.2.127.254   10.1.128.1          186
so-0/0/0         239.255.255.255 10.1.128.1          187
so-0/0/0         224.1.127.255   10.1.128.1          188
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0

user@host> clear igmp membership
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          210
so-0/0/0         239.255.255.255 10.1.128.1          210
so-0/0/0         224.1.127.255   10.1.128.1          215
so-0/0/0         224.2.127.254   10.1.128.1          216
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0

user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          210
so-0/0/0         239.255.255.255 10.1.128.1          210
so-0/0/0         224.1.127.255   10.1.128.1          215
so-0/0/0         224.2.127.254   10.1.128.1          216
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0

user@host> clear igmp membership group 239.225/16
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-1/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-2/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.1.127.255   10.1.128.1          231
so-0/0/0         224.2.127.254   10.1.128.1          233
so-0/0/0         224.2.127.253   10.1.128.1          236
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

clear igmp-snooping membership

Syntax **clear igmp-snooping membership**
<vlan *vlan-name*>

Release Information Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Clear IGMP snooping membership information.

Options **vlan** ***vlan-name***—(Optional) Name of the VLAN.

Required Privilege Level view

Related Documentation
- show igmp-snooping membership on page 364

List of Sample Output clear igmp-snooping membership on page 322

Sample Output

clear igmp-snooping membership

```
user@switch> clear igmp-snooping membership vlan employee-vlan
```

clear igmp statistics

List of Syntax
Syntax on page 323
Syntax (EX Series Switches) on page 323

Syntax
clear igmp statistics
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
clear igmp statistics
<interface *interface-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Clear Internet Group Management Protocol (IGMP) statistics.

Options
none—Clear IGMP statistics on all interfaces.

interface *interface-name*—(Optional) Clear IGMP statistics for the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- show igmp statistics on page 361

List of Sample Output
clear igmp statistics on page 323

Output Fields
See show igmp statistics for an explanation of output fields.

Sample Output

clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query           8883        459          0
V1 Membership Report          0          0          0
DVMRP                     19784      35476          0
PIM V1                    18310          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
```

```
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                  0
IGMP v3 unsupported type                             0
IGMP v3 source required for SSM                      0
IGMP v3 mode not applicable for SSM                  0

IGMP Global Statistics
Bad Length                    0
Bad Checksum                  0
Bad Receive If                0
Rx non-local               1227
```

```
user@host> clear igmp statistics
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          0          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                  0
IGMP v3 unsupported type                             0
IGMP v3 source required for SSM                      0
IGMP v3 mode not applicable for SSM                  0
IGMP Global Statistics
Bad Length                    0
Bad Checksum                  0
Bad Receive If                0
Rx non-local                  0
```

clear igmp-snooping statistics

Syntax **clear igmp-snooping statistics**

Release Information Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Clear IGMP snooping statistics.

Required Privilege Level view

Related Documentation
- show igmp-snooping statistics on page 369

List of Sample Output clear igmp-snooping statistics on page 325

Sample Output

clear igmp-snooping statistics

```
user@switch> clear igmp-snooping statistics
```

clear msdp cache

Syntax
clear msdp cache
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<peer *peer-address*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description
Clear the entries in the Multicast Source Discovery Protocol (MSDP) source-active cache.

Options
none—Clear entries in the MSDP source-active cache for all instances, logical systems, and peers.

instance *instance-name*—(Optional) Clear entries for a specific MSDP instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

peer *peer-address*—(Optional) Clear the MSDP source-active cache entries learned from a specific peer.

Required Privilege Level
clear

Related Documentation
- show msdp source-active on page 377

List of Sample Output
clear msdp cache on page 326

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear msdp cache

```
user@host> clear msdp cache
```

clear msdp statistics

Syntax
clear msdp statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<peer *peer-address*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description
Clear Multicast Source Discovery Protocol (MSDP) peer statistics.

Options
none—Clear MSDP statistics for all peers.

instance *instance-name*—(Optional) Clear statistics for the specified instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

peer *peer-address*—(Optional) Clear the statistics for the specified peer.

Required Privilege Level
clear

Related Documentation
- show msdp statistics on page 380

List of Sample Output
clear msdp statistics on page 327

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear msdp statistics

```
user@host> clear msdp statistics
```

clear multicast bandwidth-admission

Syntax

clear multicast bandwidth-admission
<group *group-address*>
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<source *source-address*>

Release Information

Command introduced in Junos OS Release 8.3.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Reapply IP multicast bandwidth admissions.

Options

none—Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.

group *group-address*—(Optional) Reapply multicast bandwidth admissions for the specified group.

inet—(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.

inet6—(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.

instance *instance-name*—(Optional) Reapply multicast bandwidth admission settings for the specified instance. If you do not specify an instance, the command applies to the master routing instance.

interface *interface-name*—(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows:

- If the interface is congested, and it was admitted previously, it is removed.
- If the interface was rejected previously, the **clear multicast bandwidth-admission** command enables the interface to be admitted as long as enough bandwidth exists on the interface.
- If you do not specify an interface, issuing the **clear multicast bandwidth-admission** command readmits any previously rejected interface for the relevant entries as long as enough bandwidth exists on the interface.

To manually reject previously admitted outbound interfaces, you must specify the interface.

source *source-address*—(Optional) Use with the **group** option to reapply multicast bandwidth admission settings for the specified (source, group) entry.

Required Privilege Level

clear

Related Documentation
- show multicast interface on page 386

List of Sample Output clear multicast bandwidth-admission on page 329

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast bandwidth-admission

```
user@host> clear multicast bandwidth-admission
```

clear multicast scope

List of Syntax
Syntax on page 330
Syntax (EX Series Switch and the QFX Series) on page 330

Syntax
clear multicast scope
<inet | inet6>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
clear multicast scope
<inet | inet6>
<interface *interface-name*>

Release Information
Command introduced in Junos OS Release 7.6.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 option introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Clear IP multicast scope statistics.

Options
none—(Same as **logical-system all**) Clear multicast scope statistics.

inet—(Optional) Clear multicast scope statistics for IPv4 family addresses.

inet6—(Optional) Clear multicast scope statistics for IPv6 family addresses.

interface *interface-name*—(Optional) Clear multicast scope statistics on a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- show multicast scope on page 407

List of Sample Output
clear multicast scope on page 330

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast scope

```
user@host> clear multicast scope
```

clear multicast sessions

List of Syntax
Syntax on page 331
Syntax (EX Series Switch and the QFX Series) on page 331

Syntax
clear multicast sessions
<logical-system (all | *logical-system-name*)>
<*regular-expression*>

Syntax (EX Series Switch and the QFX Series)
clear multicast sessions
<*regular-expression*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Clear IP multicast sessions.

Options
none—(Same as **logical-system all**) Clear multicast sessions.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

regular-expression—(Optional) Clear only multicast sessions that contain the specified regular expression.

Required Privilege Level
clear

Related Documentation
- show multicast sessions on page 409

List of Sample Output
clear multicast sessions on page 331

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast sessions

```
user@host> clear multicast sessions
```

clear multicast statistics

List of Syntax
Syntax on page 332
Syntax (EX Series Switch and the QFX Series) on page 332

Syntax

clear multicast statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

clear multicast statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Clear IP multicast statistics.

Options

none—Clear multicast statistics for all supported address families on all interfaces.

inet—(Optional) Clear multicast statistics for IPv4 family addresses.

inet6—(Optional) Clear multicast statistics for IPv6 family addresses.

instance *instance-name*—(Optional) Clear multicast statistics for the specified instance.

interface *interface-name*—(Optional) Clear multicast statistics on a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- *show multicast statistics*

List of Sample Output
clear multicast statistics on page 332

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast statistics

```
user@host> clear multicast statistics
```

clear pim join

List of Syntax
Syntax on page 333
Syntax (EX Series Switch and the QFX Series) on page 333

Syntax

clear pim join
<*group-address*>
<bidirectional | dense | sparse>
<exact>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<rp *ip-address/prefix* | source *ip-address/prefix*>
<sg | star-g>

Syntax (EX Series Switch and the QFX Series)

clear pim join
<*group-address*>
<dense | sparse>
<exact>
<inet | inet6>
<instance *instance-name*>
<rp *ip-address/prefix* | source *ip-address/prefix*>
<sg | star-g>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Multiple new filter options introduced in Junos OS Release 13.2.

Description

Clear the Protocol Independent Multicast (PIM) join and prune states.

Options

none—Clear the PIM join and prune states for all groups, family addresses, and instances.

group-address—(Optional) Clear the PIM join and prune states for a group address.

bidirectional | dense | sparse—(Optional) Clear PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.

exact—(Optional) Clear only the group that exactly matches the specified group address.

inet | inet6—(Optional) Clear the PIM entries for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear the entries for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

rp *ip-address/prefix* | source *ip-address/prefix*—(Optional) Clear the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

sg | star-g—(Optional) Clear PIM (S,G) or (*,G) entries.

Additional Information The **clear pim join** command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level clear

Related Documentation
- show pim join on page 420

List of Sample Output
clear pim join on page 334
clear pim join inet6 on page 334
clear pim join inet6 star-g on page 334

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pim join

```
user@host> clear pim join
Cleared 8 Join/Prune states
```

clear pim join inet6

```
user@host> clear pim join inet6
Cleared 4 Join/Prune states
```

clear pim join inet6 star-g

```
user@host> clear pim join inet6 star-g
Cleared 1 Join/Prune states
```

clear pim register

List of Syntax

Syntax on page 335
Syntax (EX Series Switch and the QFX Series) on page 335
Syntax (PTX Series) on page 335

Syntax

clear pim register
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

clear pim register
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>

Syntax (PTX Series)

clear pim register
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 7.6.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Clear Protocol Independent Multicast (PIM) register message counters.

Options

none—Clear PIM register message counters for all family addresses, instances, and interfaces.

inet | inet6—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear register message counters for a specific PIM-enabled routing instance.

interface *interface-name*—(Optional) Clear PIM register message counters for a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The **clear pim register** command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Related Documentation
- show pim statistics on page 448

List of Sample Output clear pim register on page 336

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pim register

```
user@host> clear pim register
```

clear pim statistics

List of Syntax
Syntax on page 337
Syntax (EX Series Switch and the QFX Series) on page 337

Syntax
clear pim statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
clear pim statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Clear Protocol Independent Multicast (PIM) statistics.

Options
none—Clear PIM statistics for all family addresses, instances, and interfaces.

inet | inet6—(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear statistics for a specific PIM-enabled routing instance.

interface *interface-name*—(Optional) Clear PIM statistics for a specific interface.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information
The **clear pim statistics** command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level
clear

Related Documentation
- show pim statistics on page 448

List of Sample Output
clear pim statistics on page 338

Output Fields
See show pim statistics for an explanation of output fields.

Sample Output

clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
Assert                        0          0          0
Graft                         0          0          0
Graft Ack                     0          0          0
Candidate RP                  0          0          0
V1 Query                   2111       4222          0
V1 Register                   0          0          0
V1 Register Stop              0          0          0
V1 Join Prune             14200      13115          0
V1 RP Reachability            0          0          0
V1 Assert                     0          0          0
V1 Graft                      0          0          0
V1 Graft Ack                  0          0          0
PIM statistics summary for all interfaces:
Unknown type                              0
V1 Unknown type                           0
Unknown Version                           0
Neighbor unknown                          0
Bad Length                                0
Bad Checksum                              0
Bad Receive If                            0
Rx Intf disabled                       2007
Rx V1 Require V2                          0
Rx Register not RP                        0
RP Filtered Source                        0
Unknown Reg Stop                          0
Rx Join/Prune no state                 1040
Rx Graft/Graft Ack no state               0
...

user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
Assert                        0          0          0
Graft                         0          0          0
Graft Ack                     0          0          0
Candidate RP                  0          0          0
V1 Query                      1          0          0
V1 Register                   0          0          0
...
```

mtrace

Syntax

mtrace *source*
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 9.5 for SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 12.3 for the PTX Series.

Description

Display trace information about an IP multicast path.

Options

source—Source hostname or address.

logical-system (*logical-system-name*)—(Optional) Perform this operation on a logical system.

routing-instance *routing-instance-name*—(Optional) Trace a particular routing instance.

Additional Information

The **mtrace** command for multicast traffic is similar to the **traceroute** command used for unicast traffic. Unlike **traceroute**, **mtrace** traces traffic backwards, from the receiver to the source.

Required Privilege Level

view

List of Sample Output

mtrace source on page 342

Output Fields

Table 9 on page 340 describes the output fields for the **mtrace** command. Output fields are listed in the approximate order in which they appear.

Table 9: mtrace Output Fields

Field Name	Field Description
Mtrace from	IP address of the receiver.
to	IP address of the source.
via group	IP address of the multicast group (if any).
Querying full reverse path	Indicates the full reverse path query has begun.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.

Table 9: mtrace Output Fields *(continued)*

Field Name	Field Description
protocol	Protocol used (for example, PIM).
Round trip time	Average round-trip time, in milliseconds (ms).
total ttl of	Time-to-live (TTL) threshold.

Sample Output

mtrace source

```
user@host> mtrace 192.1.4.2
Mtrace from 192.1.4.2 to 192.1.1.2 via group 0.0.0.0
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.1.1.2)
 -1  routerB.lab.mycompany.net (192.1.2.2)  PIM  thresh^ 1
 -2  routerC.lab.mycompany.net (192.1.3.2)  PIM  thresh^ 1
 -3  hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.
```

mtrace from-source

Syntax

mtrace from-source source *source*
<brief | detail>
<extra-hops *extra-hops*>
<group *group*>
<interval *interval*>
<loop>
<max-hops *max-hops*>
<max-queries *max-queries*>
<multicast-response | unicast-response>
<no-resolve>
<no-router-alert>
<response *response*>
<routing-instance *routing-instance-name*>
<ttl *ttl*>
<wait-time *wait-time*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display trace information about an IP multicast path from a source to this router or switch. If you specify a group address with this command, Junos OS returns additional information, such as packet rates and losses.

Options

brief | detail—(Optional) Display the specified level of output.

extra-hops *extra-hops*—(Optional) Number of hops to take after reaching a nonresponsive router. You can specify a number between **0** and **255**.

group *group*—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

interval *interval*—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10** seconds.

loop—(Optional) Loop indefinitely, displaying rate and loss statistics.

max-hops *max-hops*—(Optional) Maximum hops to trace toward the source. The range of values is **0** through **255**. The default value is **32** hops.

max-queries *max-queries*—(Optional) Maximum number of query attempts for any hop. The range of values is **1** through **32**. The default is **3**.

multicast-response—(Optional) Always request the response using multicast.

no-resolve—(Optional) Do not attempt to display addresses symbolically.

no-router-alert—(Optional) Do not use the router-alert IP option.

response *response*—(Optional) Send trace response to a host or multicast address.

routing-instance ***routing-instance-name***—(Optional) Trace a particular routing instance.

source ***source***—Source hostname or address.

ttl ***ttl***—(Optional) IP time-to-live (TTL) value. You can specify a number between **0** and **255**. Local queries to the multicast group use a value of **1**. Otherwise, the default value is **127**.

unicast-response—(Optional) Always request the response using unicast.

wait-time ***wait-time***—(Optional) Number of seconds to wait for a response. The default value is **3**.

Required Privilege Level view

List of Sample Output mtrace from-source on page 345

Output Fields Table 10 on page 344 describes the output fields for the **mtrace from-source** command. Output fields are listed in the approximate order in which they appear.

Table 10: mtrace from-source Output Fields

Field Name	Field Description
Mtrace from	IP address of the receiver.
to	IP address of the source.
via group	IP address of the multicast group (if any).
Querying full reverse path	Indicates the full reverse path query has begun.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
protocol	Protocol used (for example, PIM).
Round trip time	Average round-trip time, in milliseconds (ms).
total ttl of	Time-to-live (TTL) threshold.
source	Source address.
Response Dest	Response destination address.
Overall	Average packet rate for all traffic at each hop.

Table 10: mtrace from-source Output Fields *(continued)*

Field Name	Field Description
Packet Statistics for Traffic From	Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop.
Receiver	IP address receiving the multicast.
Query source	IP address sending the mtrace query.

Sample Output

mtrace from-source

```
user@host> mtrace from-source source 192.1.4.2 group 225.1.1.1
Mtrace from 192.1.4.2 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.1.1.2)
 -1  routerB.lab.mycompany.net (192.1.2.2)  PIM  thresh^ 1
 -2  routerC.lab.mycompany.net (192.1.3.2)  PIM  thresh^ 1
 -3  hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.

Waiting to accumulate statistics...Results after 10 seconds:

Source        Response Dest    Overall     Packet Statistics For Traffic From
192.1.4.2 192.1.1.2    Packet      192.1.4.2 To 225.1.1.1
     v       __/  rtt    2 ms     Rate       Lost/Sent = Pct  Rate
192.1.2.1
192.1.3.2  routerC.lab.mycompany.net
     v     ^      ttl    2                     0/0    = --    0 pps
192.1.4.1
192.1.2.2   routerB.lab.mycompany.net
     v      \__   ttl    3                     ?/0            0 pps
192.1.1.2   192.1.1.2
  Receiver      Query Source
```

mtrace monitor

Syntax mtrace monitor

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description Listen passively for IP multicast responses. To exit the **mtrace monitor** command, type Ctrl+c.

Options **none**—Trace the master instance.

Required Privilege Level view

List of Sample Output mtrace monitor on page 347

Output Fields Table 11 on page 346 describes the output fields for the **mtrace monitor** command. Output fields are listed in the approximate order in which they appear.

Table 11: mtrace monitor Output Fields

Field Name	Field Description
Mtrace query at	Date and time of the query.
by	Address of the host issuing the query.
resp to	Response destination.
qid	Query ID number.
packet from...to	IP address of the query source and default group destination.
from...to	IP address of the multicast source and the response address.
via group	IP address of the group to trace.
mxhop	Maximum hop setting.

Sample Output

mtrace monitor

```
user@host> mtrace monitor
Mtrace query at Oct 22 13:36:14 by 192.1.3.2, resp to 224.0.1.32, qid 74a5b8
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:17 by 192.1.3.2, resp to 224.0.1.32, qid 1d07ba
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:20 by 192.1.3.2, resp to same, qid 2fea1d
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:30 by 192.1.3.2, resp to same, qid 7c88ad
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)
```

mtrace to-gateway

Syntax

mtrace to-gateway gateway *gateway*
<brief | detail>
<extra-hops *extra-hops*>
<group *group*>
<interface *interface-name*>
<interval *interval*>
<loop>
<max-hops *max-hops*>
<max-queries *max-queries*>
<multicast-response | unicast-response>
<no-resolve>
<no-router-alert>
<response *response*>
<routing-instance *routing-instance-name*>
<ttl *ttl*>
<unicast-response>
<wait-time *wait-time*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display trace information about a multicast path from this router or switch to a gateway router or switch.

Options

gateway *gateway*—Send the trace query to a gateway multicast address.

brief | detail—(Optional) Display the specified level of output.

extra-hops *extra-hops*—(Optional) Number of hops to take after reaching a nonresponsive router or switch. You can specify a number between **0** and **255**.

group *group*—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

interface *interface-name*—(Optional) Source address for sending the trace query.

interval *interval*—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10**.

loop—(Optional) Loop indefinitely, displaying rate and loss statistics.

max-hops *max-hops*—(Optional) Maximum hops to trace toward the source. You can specify a number between **0** and **255**.. The default value is **32**.

max-queries *max-queries*—(Optional) Maximum number of query attempts for any hop. You can specify a number between **0** and **255**. The default value is **3**.

multicast-response—(Optional) Always request the response using multicast.

no-resolve—(Optional) Do not attempt to display addresses symbolically.

no-router-alert—(Optional) Do not use the router-alert IP option.

response ***response***—(Optional) Send trace response to a host or multicast address.

routing-instance ***routing-instance-name***—(Optional) Trace a particular routing instance.

ttl ***ttl***—(Optional) IP time-to-live value. You can specify a number between **0** and **225**. Local queries to the multicast group use TTL **1**. Otherwise, the default value is **127**.

unicast-response—(Optional) Always request the response using unicast.

wait-time ***wait-time***—(Optional) Number of seconds to wait for a response. The default value is **3**.

Required Privilege Level view

List of Sample Output mtrace to-gateway on page 349

Output Fields Table 12 on page 349 describes the output fields for the **mtrace to-gateway** command. Output fields are listed in the approximate order in which they appear.

Table 12: mtrace to-gateway Output Fields

Field Name	Field Description
Mtrace from	IP address of the receiver.
to	IP address of the source.
via group	IP address of the multicast group (if any).
Querying full reverse path	Indicates the full reverse path query has begun.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
protocol	Protocol used (for example, PIM).
Round trip time	Average round-trip time, in milliseconds (ms).
total ttl of	Time-to-live (TTL) threshold.

Sample Output

mtrace to-gateway

```
user@host> mtrace to-gateway gateway 192.1.3.2 group 225.1.1.1 interface 192.1.1.73 brief

Mtrace from 192.1.1.73 to 192.1.1.2 via group 225.1.1.1
```

```
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.1.1.2)
 -1  routerA.lab.mycompany.net (192.1.1.2)  PIM  thresh^ 1
 -2  routerB.lab.mycompany.net (192.1.2.2)  PIM  thresh^ 1
 -3  routerC.lab.mycompany.net (192.1.3.2)  PIM  thresh^ 1
Round trip time 2 ms; total ttl of 3 required.
```

show configuration protocols igmp

Syntax show configuration protocols igmp

Release Information Command introduced in Junos OS Release 11.3 for the QFX Series.

Description Display Internet Group Management Protocol (IGMP) information.

Required Privilege Level view

Related Documentation
- IGMP Snooping Overview on page 35
- Configuring IGMP Snooping on page 135

List of Sample Output show configuration protocols igmp on page 351

Output Fields Table 13 on page 351 describes the output fields for the **show configuration protocols igmp** command that relate to IGMP querying.

Table 13: show igmp group Output Fields

Field Name	Field Description	Level of Output
accounting	Enables notification for join and leave events.	All levels
igmp-querier	Configured source address for the IGMP querier.	All levels
interface	Name of the interface that receives IGMP membership reports.	All levels
query-interval	Interval at which the IGMP querier sends general host-query messages to solicit membership information.	All levels
query-response-interval	How long the IGMP querier waits to receive a response from a query message before sending another query.	All levels
src-address	Source address of IGMP queries.	
version	IGMP version.	All levels

Sample Output

show configuration protocols igmp

```
user@switch> show configuration protocols igmp
query-interval 150;
query-response-interval 50;
accounting;
interface vlan.43 {
 version 2;
 }
igmp-querier {
```

```
    src-address 10.0.0.2;
}
```

show igmp group

List of Syntax
Syntax on page 353
Syntax (EX Series Switch and the QFX Series) on page 353

Syntax
show igmp group
<brief | detail>
<*group-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show igmp group
<brief | detail>
<*group-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display Internet Group Management Protocol (IGMP) group membership information.

Options
none—Display standard information about membership for all IGMP groups.

brief | detail—(Optional) Display the specified level of output.

group-name—(Optional) Display group membership for the specified IP address only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear igmp membership on page 319

List of Sample Output
show igmp group (Include Mode) on page 354
show igmp group (Exclude Mode) on page 355
show igmp group brief on page 355
show igmp group detail on page 355

Output Fields
Table 13 on page 351 describes the output fields for the **show igmp group** command. Output fields are listed in the approximate order in which they appear.

Table 14: show igmp group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the IGMP membership report. A name of **local** indicates that the local routing device joined the group itself.	All levels
Group	Group address.	All levels
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.	All levels

Table 14: show igmp group Output Fields *(continued)*

Field Name	Field Description	Level of Output
Source	Source address.	All levels
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	**detail**
Last reported by	Address of the host that last reported membership in this group.	All levels
Timeout	Time remaining until the group membership is removed.	**brief none**
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	**detail**
Type	Type of group membership: • **Dynamic**—Host reported the membership. • **Static**—Membership is configured.	All levels

Sample Output

show igmp group (Include Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.2
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.3
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.4
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
    Group: 232.1.1.2
        Group mode: Include
        Source: 10.0.0.4
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 224.0.0.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

```
    Group: 224.0.0.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show igmp group (Exclude Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 224.0.0.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 224.0.0.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

show igmp group detail

```
user@host> show igmp group detail
Interface: t1-0/1/0.0
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.2
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.3
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.4
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
    Group: 232.1.1.2
        Group mode: Include
        Source: 10.0.0.4
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 224.0.0.2
        Group mode: Exclude
```

```
    Source: 0.0.0.0
    Source timeout: 0
    Last reported by: Local
    Group timeout:       0 Type: Dynamic
Group: 224.0.0.22
    Group mode: Exclude
    Source: 0.0.0.0
    Source timeout: 0
    Last reported by: Local
    Group timeout:       0 Type: Dynamic
```

show igmp interface

List of Syntax
Syntax on page 357
Syntax (EX Series Switches and the QFX Series) on page 357

Syntax
show igmp interface
<brief | detail>
<*interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and the QFX Series)
show igmp interface
<brief | detail>
<*interface-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.

Options
none—Display standard information about all IGMP-enabled interfaces.

brief | detail—(Optional) Display the specified level of output.

interface-name—(Optional) Display information about the specified IGMP-enabled interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear igmp membership on page 319

List of Sample Output
show igmp interface on page 359
show igmp interface brief on page 360
show igmp interface detail on page 360
show igmp interface <interface-name> on page 360

Output Fields
Table 15 on page 357 describes the output fields for the **show igmp interface** command. Output fields are listed in the approximate order in which they appear.

Table 15: show igmp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the routing device that has been elected to send membership queries.	All levels

Table 15: show igmp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
State	State of the interface: **Up** or **Down**.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.	All levels
Timeout	How long until the IGMP querier is declared to be unreachable, in seconds.	All levels
Version	IGMP version being used on the interface: **1** , **2** , or **3**.	All levels
Groups	Number of groups on the interface.	All levels
Group limit	Maximum number of groups allowed on the interface. Any joins requested after the limit is reached are rejected.	All levels
Group threshold	Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.	All levels
Group log-interval	Time (in seconds) between consecutive log messages.	All levels
Immediate Leave	State of the immediate leave option: • **On**—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface. • **Off**—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.	All levels
Promiscuous Mode	State of the promiscuous mode option: • **On**—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces. • **Off**—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces.	All levels
Passive	State of the passive mode option: • **On**—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves. • **Off**—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves. The **passive** statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the **on** state declaration: • **send-general-query**—The interface sends general queries. • **send-group-query**—The interface sends group-specific and group-source-specific queries. • **allow-receive**—The interface receives control traffic.	All levels

Table 15: show igmp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
OIF map	Name of the OIF map (if configured) associated with the interface.	All levels
SSM map	Name of the source-specific multicast (SSM) map (if configured) used on the interface.	All levels
Configured Parameters	Information configured by the user: • **IGMP Query Interval**—Interval (in seconds) at which this router sends membership queries when it is the querier. • **IGMP Query Response Interval**—Time (in seconds) that the router waits for a report in response to a general query. • **IGMP Last Member Query Interval**—Time (in seconds) that the router waits for a report in response to a group-specific query. • **IGMP Robustness Count**—Number of times the router retries a query.	All levels
Derived Parameters	Derived information: • **IGMP Membership Timeout**—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed. • **IGMP Other Querier Present Timeout**—Time (in seconds) that the router waits for the IGMP querier to send a query.	All levels

Sample Output

show igmp interface

```
user@host> show igmp interface
Interface: at-0/3/1.0
    Querier: 10.111.30.1
    State:         Up Timeout:    None Version:  2 Groups:      4
    SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
    Querier: 10.111.10.1
    State:         Up Timeout:    None Version:  2 Groups:      2
    SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
    Querier: 10.111.20.1
    State:         Up Timeout:    None Version:  2 Groups:      4
    SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see show igmp interface on page 359.

show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see show igmp interface on page 359.

show igmp interface <interface-name>

```
user@host# show igmp interface ge-3/2/0.0
Interface: ge-3/2/0.0
  Querier: 20.1.1.1
  State: Up Timeout:    None Version:  3 Groups:      1
  Group limit: 8
  Group threshold: 60
  Group log-interval: 10
  Immediate leave:  Off
  Promiscuous mode:  Off
```

show igmp statistics

List of Syntax
Syntax on page 361
Syntax (EX Series Switch and the QFX Series) on page 361

Syntax
show igmp statistics
<brief | detail>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show igmp statistics
<brief | detail>
<interface *interface-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display Internet Group Management Protocol (IGMP) statistics.

Options
none—Display IGMP statistics for all interfaces.

brief | detail—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display IGMP statistics about the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear igmp statistics on page 323

List of Sample Output
show igmp statistics on page 362
show igmp statistics interface on page 363

Output Fields
Table 16 on page 361 describes the output fields for the **show igmp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 16: show igmp statistics Output Fields

Field Name	Field Description
IGMP packet statistics	Heading for IGMP packet statistics for all interfaces or for the specified interface name.

Table 16: show igmp statistics Output Fields *(continued)*

Field Name	Field Description
IGMP Message type	Summary of IGMP statistics: • **Membership Query**—Number of membership queries sent and received. • **V1 Membership Report**—Number of version 1 membership reports sent and received. • **DVMRP**—Number of DVMRP messages sent or received. • **PIM V1**—Number of PIM version 1 messages sent or received. • **Cisco Trace**—Number of Cisco trace messages sent or received. • **V2 Membership Report**—Number of version 2 membership reports sent or received. • **Group Leave**—Number of group leave messages sent or received. • **Mtrace Response**—Number of Mtrace response messages sent or received. • **Mtrace Request**—Number of Mtrace request messages sent or received. • **Domain Wide Report**—Number of domain-wide reports sent or received. • **V3 Membership Report**—Number of version 3 membership reports sent or received. • **Other Unknown types**—Number of unknown message types received. • **IGMP v3 unsupported type**—Number of messages received with unknown and unsupported IGMP version 3 message types. • **IGMP v3 source required for SSM**—Number of IGMP version 3 messages received that contained no source. • **IGMP v3 mode not applicable for SSM**—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).
Received	Number of messages received.
Sent	Number of messages sent.
Rx errors	Number of received packets that contained errors.
IGMP Global Statistics	Summary of IGMP statistics for all interfaces. • **Bad Length**—Number of messages received with length errors so severe that further classification could not occur. • **Bad Checksum**—Number of messages received with a bad IP checksum. No further classification was performed. • **Bad Receive If**—Number of messages received on an interface not enabled for IGMP. • **Rx non-local**—Number of messages received from senders that are not local. • **Timed out**—Number of groups that timed out as a result of not receiving an explicit leave message. • **Rejected Report**—Number of reports dropped because of the IGMP group policy. • **Total Interfaces**—Number of interfaces configured to support IGMP.

Sample Output

show igmp statistics

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query           8883        459          0
V1 Membership Report          0          0          0
```

```
DVMRP                          0          0          0
PIM V1                         0          0          0
Cisco Trace                    0          0          0
V2 Membership Report           0          0          0
Group Leave                    0          0          0
Mtrace Response                0          0          0
Mtrace Request                 0          0          0
Domain Wide Report             0          0          0
V3 Membership Report           0          0          0
Other Unknown types                                  0
IGMP v3 unsupported type                             0
IGMP v3 source required for SSM                      0
IGMP v3 mode not applicable for SSM                  0

IGMP Global Statistics
Bad Length                     0
Bad Checksum                   0
Bad Receive If                 0
Rx non-local                1227
Timed out                      0
Rejected Report                0
Total Interfaces               2
```

show igmp statistics interface

```
user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type      Received       Sent  Rx errors
Membership Query              0        230          0
V1 Membership Report          0          0          0
```

show igmp-snooping membership

Syntax

show igmp-snooping membership
<**brief** | **detail**>
<**interface** *interface-name*>
<**vlan** *vlan-id* | *vlan-name*>

Release Information

Command introduced in Junos OS Release 11.1 for the QFX Series.
IGMPv3 output introduced in Junos OS Release 12.1 for the QFX Series.

Description

Display IGMP snooping membership information.

Options

none—Display general parameters.

brief | detail—(Optional) Display the specified level of output.

interface ***interface-name***—(Optional) Display IGMP snooping information for the specified interface.

vlan ***vlan-id*** **|** ***vlan-name***—(Optional) Display IGMP snooping information for the specified VLAN.

Required Privilege Level

view

Related Documentation

- Monitoring IGMP Snooping on page 315
- Configuring IGMP Snooping on page 135
- show igmp-snooping route on page 367
- show igmp-snooping statistics on page 369
- show igmp-snooping vlans on page 371

List of Sample Output

show igmp-snooping membership on page 365
show igmp-snooping membership detail on page 366

Output Fields

Table 17 on page 364 lists the output fields for the **show igmp-snooping membership** command. Output fields are listed in the approximate order in which they appear.

Table 17: show igmp-snooping membership Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All
Interfaces	Interfaces assigned to the VLAN.	All
Tag	Numerical identifier of the VLAN.	**detail**

Table 17: show igmp-snooping membership Output Fields *(continued)*

Field Name	Field Description	Level of Output
Router interfaces	Names of multicast router interfaces.	**detail**
• **static** or **dynamic**	Whether the multicast router interface is **static** or **dynamic**.	**detail**
• **Uptime**	For static interfaces, length of time since the interface was configured as a multicast router interface; for dynamic interfaces, length of time since the first query was received on the interface.	**detail**
• **timeout**	Query timeout in seconds.	**detail**
Group	IP multicast address of the multicast group.	**detail**
Receiver count	Number of interfaces that have membership in a multicast group.	**detail**
Flags	IGMP version of the host sending a join message.	**detail**
Uptime	Length of time a multicast group has been active on the interface.	**detail**
timeout	Time (in seconds) left until the entry for the multicast group is removed.	All
Last reporter	Last host to report membership for the multicast group.	**detail**
Include source	Source addresses from which multicast streams are allowed based on IGMPv3 reports.	**detail**

Sample Output

show igmp-snooping membership

```
user@switch> show igmp-snooping membership
VLAN: v1
    224.1.1.1      *              258 secs
        Interfaces: ge-0/0/0.0
    224.1.1.3      *              258 secs
        Interfaces: ge-0/0/0.0
    224.1.1.5      *              258 secs
        Interfaces: ge-0/0/0.0
    224.1.1.7      *              258 secs
```

```
        Interfaces: ge-0/0/0.0
    224.1.1.9      *                258 secs
        Interfaces: ge-0/0/0.0
    224.1.1.11     *                258 secs
        Interfaces: ge-0/0/0.0
```

show igmp-snooping membership detail

```
user@switch> show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
   Group: 225.0.0.2
   Receiver count: 1, Flags: <V3-hosts>
     ge-0/0/15.0 Uptime: 00:00:11 timeout: 248 Last reporter: 10.2.10.16
       Include source: 1.2.1.1, 1.3.1.1
VLAN: v44 Tag: 44 (Index: 5)
   Group: 225.0.0.1
   Receiver count: 1, Flags: <V2-hosts>
     ge-0/0/21.0 Uptime: 00:00:02 timeout: 257
VLAN: v110 Tag: 110 (Index: 4)
 Router interfaces:
         ge-0/0/3.0 static Uptime: 00:08:45
         ge-0/0/2.0 static Uptime: 00:08:45
         ge-0/0/4.0 dynamic Uptime: 00:16:41 timeout: 254
     Group: 225.0.0.3
     Receiver count: 1, Flags: <V3-hosts>
         ge-0/0/5.0 Uptime: 00:00:19 timeout: 259
     Group: 225.1.1.1
     Receiver count: 1, Flags: <V2-hosts>
         ge-0/0/5.0 Uptime: 00:22:43 timeout: 96
     Group: 225.2.2.2
     Receiver count: 1, Flags: <V2-hosts Static>
         ge-0/0/5.0 Uptime: 00:23:13
```

show igmp-snooping route

Syntax

show igmp-snooping route
<**brief** | **detail**>
<**ethernet-switching** <**brief** | **detail** | **vlan** (*vlan-id* | *vlan-name*)>>
<**inet** <**brief** | **detail** | **vlan** *vlan-name*>>
<**vlan** *vlan-name*>

Release Information

Command introduced in Junos OS Release 11.1 for the QFX Series.

Description

Display IGMP snooping route information.

Options

none—Display general parameters.

brief | detail—(Optional) Display the specified level of output.

ethernet-switching—(Optional) Display Ethernet switching information.

inet—(Optional) Display **inet** information.

vlan *vlan-name*—(Optional) Display route information for the specified VLAN.

Required Privilege Level

view

Related Documentation

- Monitoring IGMP Snooping on page 315
- Configuring IGMP Snooping on page 135
- show igmp-snooping statistics on page 369
- show igmp-snooping vlans on page 371

List of Sample Output

show igmp-snooping route on page 368
show igmp-snooping route vlan v1 on page 368

Output Fields

Table 18 on page 367 lists the output fields for the **show igmp-snooping route** command. Output fields are listed in the approximate order in which they appear.

Table 18: show igmp-snooping route Output Fields

Field Name	Field Description
Table	(For internal use only. Value is always 0.)
VLAN	Name of the VLAN.
Group	Multicast group address.
Interfaces	Interfaces on which IGMP packets were snooped.
Next-hop	ID associated with the next-hop device.

Sample Output

show igmp-snooping route

```
user@switch> show igmp-snooping route
VLAN             Group          Next-hop
V11              224.1.1.1, *        533
        Interfaces: ge-0/0/13.0, ge-0/0/1.0
VLAN             Group          Next-hop
v12              224.1.1.3, *        534
        Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

show igmp-snooping route vlan v1

```
user@switch> show igmp-snooping route vlan v1
Table: 0
VLAN             Group          Next-hop
v1               224.1.1.1, *       1266
        Interfaces: ge-0/0/0.0
v1               224.1.1.3, *       1266
        Interfaces: ge-0/0/0.0
v1               224.1.1.5, *       1266
        Interfaces: ge-0/0/0.0
v1               224.1.1.7, *       1266
        Interfaces: ge-0/0/0.0
v1               224.1.1.9, *       1266
        Interfaces: ge-0/0/0.0
v1               224.1.1.11, *      1266
        Interfaces: ge-0/0/0.0
```

show igmp-snooping statistics

Syntax **show igmp-snooping statistics**

Release Information Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display IGMP snooping statistics.

Required Privilege Level view

Related Documentation

- Monitoring IGMP Snooping on page 315
- Configuring IGMP Snooping on page 135
- show igmp-snooping route on page 367
- show igmp-snooping vlans on page 371

List of Sample Output show igmp-snooping statistics on page 370

Output Fields Table 19 on page 369 lists the output fields for the **show igmp-snooping statistics** command. Output fields are listed in the approximate order in which they appear.

Table 19: show igmp-snooping statistics Output Fields

Field Name	Field Description
Bad length	IGMP packet has illegal or bad length.
Bad checksum	IGMP or IP checksum is incorrect.
Invalid interface	Packet was received through an invalid interface.
Not local	Number of packets received from senders that are not local.
Receive unknown	Unknown IGMP type.
Timed out	Number of timeouts for all multicast groups.
IGMP Type	Type of IGMP message (**Queries**, **Reports**, **Leaves**, or **Other**).
Received	Number of IGMP packets received.
Transmitted	Number of IGMP packets transmitted.
Recv Errors	Number of general receive errors.

Sample Output

show igmp-snooping statistics

```
user@switch> show igmp-snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 58

    IGMP Type           Received      Transmitted    Recv Errors
    Queries:               74295                0              0
    Reports:            18148423                0       16333523
    Leaves:                    0                0              0
    Other:                     0                0              0
```

show igmp-snooping vlans

Syntax
show igmp-snooping vlans
<**brief** | **detail**>
<**vlan** ***vlan-id*** | *vlan-name*>

Release Information Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display IGMP snooping VLAN information.

Options
none—Display general parameters.

brief | detail—(Optional) Display the specified level of output.

vlan *vlan-id* | vlan *vlan-number*—(Optional) Display VLAN information for the specified VLAN.

Required Privilege Level view

Related Documentation
- Monitoring IGMP Snooping on page 315
- Configuring IGMP Snooping on page 135
- show igmp-snooping route on page 367
- show igmp-snooping statistics on page 369

List of Sample Output
show igmp-snooping vlans on page 372
show igmp-snooping vlans vlan on page 372
show igmp-snooping vlans vlan detail on page 372

Output Fields Table 20 on page 371 lists the output fields for the **show igmp-snooping vlans** command. Output fields are listed in the approximate order in which they appear.

Table 20: show igmp-snooping vlans Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All levels
IGMP-L2-Querier	Source address for IGMP snooping queries (if switch is an IGMP querier)	All levels
Interfaces	Number of interfaces in the VLAN.	All levels
Groups	Number of groups in the VLAN.	All levels
MRouters	Number of multicast routers associated with the VLAN.	All levels
Receivers	Number of host receivers in the VLAN.	All levels

Table 20: show igmp-snooping vlans Output Fields *(continued)*

Field Name	Field Description	Level of Output
Tag	Numerical identifier of the VLAN.	**detail**
tagged | untagged	Interface participates in a tagged (802.1Q) or untagged (native) VLAN.	**detail**
vlan-interface	Internal VLAN interface identifier.	**detail**
Membership timeout	Membership timeout value.	**detail**
Querier timeout	Timeout value for interfaces dynamically marked as router or switch interfaces (interfaces that receive queries). When the querier timeout is reached, the switch marks the interface as a host interface.	**detail**
Interface	Name of the interface.	**detail**
Reporters	Number of dynamic groups on an interface.	**detail**

Sample Output

show igmp-snooping vlans

```
user@switch> show igmp-snooping vlans
VLAN     Interfaces Groups MRouters Receivers
default     0          0      0         0
v1          11         50     0         0
v10         1          0      0         0
v11         1          0      0         0
v180        3          0      1         0
v181        3          0      0         0
v182        3          0      0         0
```

show igmp-snooping vlans vlan

```
user@switch> show igmp-snooping vlans vlan v10
user@switch> show igmp-snooping vlans vlan v10
VLAN       Interfaces Groups MRouters Receivers
v10        1          0      0        0
```

show igmp-snooping vlans vlan detail

```
user@switch> show igmp-snooping vlans vlan v10 detail
VLAN: v10, Tag: 10, vlan-interface: vlan.10
        Interface: ge-0/0/10.0, tagged, Groups: 0
IGMP-L2-Querier: Stopped, SourceAddress: 10.10.1.2
```

show msdp

Syntax

show msdp
<brief | detail>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<peer *peer-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description

Display Multicast Source Discovery Protocol (MSDP) information.

Options

none—Display standard MSDP information for all routing instances.

brief | detail—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

peer *peer-address*—(Optional) Display information about the specified peer only,

Required Privilege Level

view

Related Documentation

- show msdp source on page 375
- show msdp source-active on page 377
- show msdp statistics on page 380

List of Sample Output

show msdp on page 374
show msdp brief on page 374
show msdp detail on page 374

Output Fields

Table 21 on page 373 describes the output fields for the **show msdp** command. Output fields are listed in the approximate order in which they appear.

Table 21: show msdp Output Fields

Field Name	Field Description	Level of Output
Peer address	IP address of the peer.	All levels
Local address	Local address of the peer.	All levels
State	Status of the MSDP connection: **Listen**, **Established**, or **Inactive**.	All levels
Last up/down	Time at which the most recent peer-state change occurred.	All levels

Table 21: show msdp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Peer-Group	Peer group name.	All levels
SA Count	Number of source-active cache entries advertised by each peer that were accepted, compared to the number that were received, in the format *number-accepted/number-received*.	All levels
Peer Connect Retries	Number of peer connection retries.	**detail**
State timer expires	Number of seconds before another message is sent to a peer.	**detail**
Peer Times out	Number of seconds to wait for a response from the peer before the peer is declared unavailable.	**detail**
SA accepted	Number of entries in the source-active cache accepted from the peer.	**detail**
SA received	Number of entries in the source-active cache received by the peer.	**detail**

Sample Output

show msdp

```
user@host> show msdp
Peer address     Local address   State       Last up/down Peer-Group SA Count
198.32.8.193     198.32.8.195    Established  5d 19:25:44 North23    120/150
198.32.8.194     198.32.8.195    Established  3d 19:27:27 North23    300/345
198.32.8.196     198.32.8.195    Established  5d 19:39:36 North23    10/13
198.32.8.197     198.32.8.195    Established  5d 19:32:27 North23    5/6
198.32.8.198     198.32.8.195    Established  3d 19:33:04 North23    2305/3000
```

show msdp brief

The output for the **show msdp brief** command is identical to that for the **show msdp** command. For sample output, see show msdp on page 374.

show msdp detail

```
user@host> show msdp detail
Peer: 10.255.70.15
Local address: 10.255.70.19
State: Established
Peer Connect Retries: 0
State timer expires: 22
Peer Times out: 49
SA accepted: 0
SA received: 0
```

show msdp source

Syntax

show msdp source
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*source-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description

Display multicast sources learned from Multicast Source Discovery Protocol (MSDP).

Options

none—Display standard MSDP source information for all routing instances.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

source-address—(Optional) IP address and optional prefix length. Display information for the specified source address only.

Required Privilege Level

view

Related Documentation

- show msdp on page 373
- show msdp source-active on page 377
- show msdp statistics on page 380

List of Sample Output

show msdp source on page 376

Output Fields Table 22 on page 376 describes the output fields for the **show msdp source** command. Output fields are listed in the approximate order in which they appear.

Table 22: show msdp source Output Fields

Field Name	Field Description
Source address	IP address of the source.
/Len	Length of the prefix for this IP address.
Type	Discovery method for this multicast source: • **Configured**—Source-active limit explicitly configured for this source. • **Dynamic**—Source-active limit established when this source was discovered.
Maximum	Source-active limit applied to this source.
Threshold	Source-active threshold applied to this source.
Exceeded	Number of source-active messages received from this source exceeding the established maximum.

Sample Output

show msdp source

```
user@host> show msdp source
Source address /Len   Type       Maximum   Threshold    Exceeded
0.0.0.0        /0     Configured        5       none           0
10.1.0.0       /16    Configured      500       none           0
10.1.1.1       /32    Configured    10000       none           0
10.1.1.2       /32    Dynamic        6936       none           0
10.1.5.5       /32    Dynamic         500       none         123
10.2.1.1       /32    Dynamic           2       none           0
```

show msdp source-active

Syntax

show msdp source-active
<brief | detail>
<group *group*>
<instance *instance-name*>
<local>
<logical-system (all | *logical-system-name*)>
<originator *originator*>
<peer *peer-address*>
<source *source-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description

Display the Multicast Source Discovery Protocol (MSDP) source-active cache.

Options

none—Display standard MSDP source-active cache information for all routing instances.

brief | detail—(Optional) Display the specified level of output.

group *group*—(Optional) Display source-active cache information for the specified group.

instance *instance-name*—(Optional) Display information for the specified instance.

local—(Optional) Display all source-active caches originated by this router.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

originator *originator*—(Optional) Display information about the peer that originated the source-active cache entries.

peer *peer-address*—(Optional) Display the source-active cache of the specified peer.

source *source-address*—(Optional) Display the source-active cache of the specified source.

Required Privilege Level

view

Related Documentation

- show msdp on page 373
- show msdp source on page 375
- show msdp statistics on page 380

List of Sample Output

Output Fields

Table 23 on page 378 describes the output fields for the **show msdp source-active** command. Output fields are listed in the approximate order in which they appear.

Table 23: show msdp source-active Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.
Group address	Multicast address of the group.
Source address	IP address of the source.
Peer address	IP address of the peer.
Originator	Router ID configured on the source of the rendezvous point (RP) that originated the message, or the loopback address when the router ID is not configured.
Flags	Flags: **Accept**, **Reject**, or **Filtered**.

Sample Output

show msdp source-active

```
user@host> show msdp source-active
Group address   Source address  Peer address    Originator      Flags
230.0.0.0       192.168.195.46  local           10.255.14.30    Accept
230.0.0.1       192.168.195.46  local           10.255.14.30    Accept
230.0.0.2       192.168.195.46  local           10.255.14.30    Accept
230.0.0.3       192.168.195.46  local           10.255.14.30    Accept
230.0.0.4       192.168.195.46  local           10.255.14.30    Accept
```

show msdp source-active brief

The output for the **show msdp source-active brief** command is identical to that for the **show msdp source-active** command. For sample output, see show msdp source-active on page 378.

show msdp source-active detail

The output for the **show msdp source-active detail** command is identical to that for the **show msdp source-active** command. For sample output, see show msdp source-active on page 378.

show msdp source-active source

```
user@host> show msdp source-active source 192.168.215.246
Global active source limit exceeded: 0
Global active source limit maximum: 25000
Global active source limit threshold: 24000
Global active source limit log-warning: 100
Global active source limit log interval: 0

Group address   Source address  Peer address    Originator      Flags
226.2.2.1       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.3       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.4       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.5       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.7       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.10      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.11      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.13      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.14      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.15      192.168.215.246 10.255.182.140  10.255.182.140  Accept
```

show msdp statistics

Syntax

show msdp statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<peer *peer-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description

Display statistics about Multicast Source Discovery Protocol (MSDP) peers.

Options

none—Display statistics about all MSDP peers for all routing instances.

instance *instance-name*—(Optional) Display statistics about a specific MSDP instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

peer *peer-address*—(Optional) Display statistics about a particular MSDP peer.

Required Privilege Level

view

Related Documentation

- clear msdp statistics on page 327

List of Sample Output

show msdp statistics on page 382
show msdp statistics peer on page 382

Output Fields

Table 24 on page 380 describes the output fields for the **show msdp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 24: show msdp statistics Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.
Peer	Address of peer.

Table 24: show msdp statistics Output Fields *(continued)*

Field Name	Field Description
Last State Change	How long ago the peer state changed.
Last message received from the peer	How long ago the last message was received from the peer.
RPF Failures	Number of reverse path forwarding (RPF) failures.
Remote Closes	Number of times the remote peer closed.
Peer Timeouts	Number of peer timeouts.
SA messages sent	Number of source-active messages sent.
SA messages received	Number of source-active messages received.
SA request messages sent	Number of source-active request messages sent.
SA request messages received	Number of source-active request messages received.
SA response messages sent	Number of source-active response messages sent.
SA response messages received	Number of source-active response messages received.
Active source exceeded	Number of times this peer has exceeded configured source-active limits.
Active source Maximum	Configured number of active source messages accepted by this peer.
Active source threshold	Configured threshold on this peer for applying random early discard (RED) to drop some but not all MSDP active source messages.
Active source log-warning	Configured threshold on this peer at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Active source log-interval	Time (in seconds) between consecutive log messages on this peer.
Keepalive messages sent	Number of keepalive messages sent.
Keepalive messages received	Number of keepalive messages received.
Unknown messages received	Number of unknown messages received.

Table 24: show msdp statistics Output Fields *(continued)*

Field Name	Field Description
Error messages received	Number of error messages received.

Sample Output

show msdp statistics

```
user@host> show msdp statistics
Global active source limit exceeded: 0
Global active source limit maximum: 10
Global active source limit threshold: 8
Global active source limit log-warning: 60
Global active source limit log interval: 60

Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 0
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval 120
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0
```

show msdp statistics peer

```
user@host> show msdp statistics peer 10.255.182.140
Peer: 10.255.182.140
    Last State Change: 8:19:23 (00:01:08)
    Last message received from peer: 8:20:05 (00:00:26)
    RPF Failures: 0
    Remote Closes: 0
    Peer Timeouts: 0
    SA messages sent: 17
    SA messages received: 16
    SA request messages sent: 0
    SA request messages received: 0
    SA response messages sent: 0
    SA response messages received: 0
    Active source exceeded: 20
    Active source Maximum: 10
    Active source threshold: 8
    Active source log-warning: 60
    Active source log-interval: 120
    Keepalive messages sent: 0
```

```
Keepalive messages received: 0
Unknown messages received: 0
Error messages received: 0
```

show multicast flow-map

List of Syntax
Syntax on page 384
Syntax (EX Series Switch and the QFX Series) on page 384

Syntax
show multicast flow-map
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show multicast flow-map
<brief | detail>

Release Information
Command introduced in Junos OS Release 8.2.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display configuration information about IP multicast flow maps.

Options
none—Display configuration information about IP multicast flow maps on all systems.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast flow-map on page 385
show multicast flow-map detail on page 385

Output Fields
Table 25 on page 384 describes the output fields for the **show multicast flow-map** command. Output fields are listed in the approximate order in which they appear.

Table 25: show multicast flow-map Output Fields

Field Name	Field Description	Levels of Output
Name	Name of the flow map.	All levels
Policy	Name of the policy associated with the flow map.	All levels
Cache-timeout	Cache timeout value assigned to the flow map.	All levels
Bandwidth	Bandwidth setting associated with the flow map.	All levels
Adaptive	Whether or not adaptive mode is enabled for the flow map.	none
Flow-map	Name of the flow map.	**detail**

Table 25: show multicast flow-map Output Fields *(continued)*

Field Name	Field Description	Levels of Output
Adaptive Bandwidth	Whether or not adaptive mode is enabled for the flow map.	**detail**
Redundant Sources	Redundant sources defined for the same destination group.	**detail**

Sample Output

show multicast flow-map

```
user@host> show multicast flow-map
Instance: master
Name            Policy                  Cache timeout        Bandwidth Adaptive
map2            policy2                         never          2000000 no
map1            policy1                    60 seconds          2000000 no
```

Sample Output

show multicast flow-map detail

```
user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
    Policy:             policy1
    Cache Timeout:      600 seconds
    Bandwidth:          2000000
    Adaptive Bandwidth: yes
    Redundant Sources:  11.11.11.11
    Redundant Sources:  11.11.11.12
    Redundant Sources:  11.11.11.13
```

show multicast interface

List of Syntax
Syntax on page 386
Syntax (EX Series Switch and the QFX Series) on page 386

Syntax
show multicast interface
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show multicast interface

Release Information
Command introduced in Junos OS Release 8.3.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display bandwidth information about IP multicast interfaces.

Options
none—Display all interfaces that have multicast configured.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast interface on page 387

Output Fields
Table 26 on page 386 describes the output fields for the **show multicast interface** command. Output fields are listed in the approximate order in which they appear.

Table 26: show multicast interface Output Fields

Field Name	Field Description
Interface	Name of the multicast interface.
Maximum bandwidth (bps)	Maximum bandwidth setting, in bits per second, for this interface.
Remaining bandwidth (bps)	Amount of bandwidth, in bits per second, remaining on the interface.
Mapped bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface. NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. This field does not appear in the output when the no QoS adjustment feature is disabled.

Table 26: show multicast interface Output Fields *(continued)*

Field Name	Field Description
Local bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface. NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. This field does not appear in the output when the no QoS adjustment feature is disabled.
Reverse OIF mapping	State of the reverse OIF mapping feature (**on** or **off**). NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
Reverse OIF mapping no QoS adjustment	State of the no QoS adjustment feature (**on** or **off**) for interfaces that are using reverse OIF mapping. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
Leave timer	Amount of time a mapped interface remains active after the last mapping ends. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
No QoS adjustment	State (**on**) of the no QoS adjustment feature when this feature is enabled. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.

Sample Output

show multicast interface

```
user@host> show multicast interface
Interface               Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3                              10000000                          0
fe-0/0/3.210                          10000000                   -2000000
fe-0/0/3.220                         100000000                  100000000
fe-0/0/3.230                          20000000                   18000000
fe-0/0/2.200                         100000000                  100000000
```

show multicast mrinfo

Syntax show multicast mrinfo
<host>

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description Display configuration information about IP multicast networks, including neighboring multicast router addresses.

Options **none**—Display configuration information about all multicast networks.

host—(Optional) Display configuration information about a particular host. Replace *host* with a hostname or IP address.

Required Privilege Level view

List of Sample Output show multicast mrinfo on page 389

Output Fields Table 27 on page 388 describes the output fields for the **show multicast mrinfo** command. Output fields are listed in the approximate order in which they appear.

Table 27: show multicast mrinfo Output Fields

Field Name	Field Description
source-address	Query address, hostname (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor.
ip-address-1—>ip-address-2	Queried router interface address and directly attached neighbor interface address, respectively.
(*name* or *ip-address*)	Name or IP address of neighbor.

[*metric/threshold/type/flags*]	Neighbor's multicast profile:
	• *metric*—Always has a value of **1**, because **mrinfo** queries the directly connected interfaces of a device.
	• *threshold*—Multicast threshold time-to-live (TTL). The range of values is **0** through **255**.
	• *type*—Multicast connection type: **pim** or **tunnel**.
	• *flags*—Flags for this route:
	  ◦ **querier**—Queried router is the designated router for the neighboring session.
	  ◦ **leaf**—Link is a leaf in the multicast network.
	  ◦ **down**—Link status indicator.

Sample Output

show multicast mrinfo

```
user@host> show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
  192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
  0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```

show multicast next-hops

List of Syntax
Syntax on page 390
Syntax (EX Series Switch and the QFX Series) on page 390

Syntax
show multicast next-hops
<brief | detail>
<*identifier-number*>
<inet | inet6>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show multicast next-hops
<brief | detail>
<*identifier-number*>
<inet | inet6>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 option introduced in Junos OS Release 10.0 for EX Series switches.
detail option display of next-hop ID number introduced in Junos OS Release 11.1 for M Series and T Series routers and EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.

Description
Display the entries in the IP multicast next-hop table.

Options
none—Display standard information about all entries in the multicast next-hop table for all supported address families.

brief | detail—(Optional) Display the specified level of output.

When you include the **detail** option on M Series and T Series routers and EX Series switches, the downstream interface name includes the next-hop ID number in parentheses, in the form **fe-0/1/2.0-(1048574)** where **1048574** is the next-hop ID number.

identifier-number—(Optional) Show a particular next hop by ID number. The range of values is **1** through **65,535**.

inet | inet6—(Optional) Display entries for IPv4 or IPv6 family addresses, respectively.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast next-hops on page 391
show multicast next-hops (Bidirectional PIM on page 391
show multicast next-hops brief on page 392
show multicast next-hops detail on page 392

Output Fields Table 28 on page 391 describes the output fields for the **show multicast next-hops** command. Output fields are listed in the approximate order in which they appear.

Table 28: show multicast next-hops Output Fields

Field Name	Field Description
Family	Protocol family (such as **INET**).
ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine.
Refcount	Number of cache entries that are using this next hop.
KRefcount	Kernel reference count for the next hop.
Downstream interface	Interface names associated with each multicast next-hop ID.
Incoming interface list	List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.

Sample Output

show multicast next-hops

```
user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefcount Downstream interface
262142          4          2 so-1/0/0.0
262143          2          1 mt-1/1/0.49152
262148          2          1 mt-1/1/0.32769
```

show multicast next-hops (Bidirectional PIM

```
user@host> show multicast next-hops
Family: INET
ID          Refcount KRefcount Downstream interface
2097151            8         4 ge-0/0/1.0

Family: INET6
ID          Refcount KRefcount Downstream interface
2097157            2         1 ge-0/0/1.0

Family: Incoming interface list
ID          Refcount KRefcount Downstream interface
513                5         2 lo0.0
                               ge-0/0/1.0
514                5         2 lo0.0
                               ge-0/0/1.0
                               xe-4/1/0.0
515                3         1 lo0.0
                               ge-0/0/1.0
                               xe-4/1/0.0
544                1         0 lo0.0
                               xe-4/1/0.0
```

show multicast next-hops brief

The output for the **show multicast next-hops brief** command is identical to that for the **show multicast next-hops** command. For sample output, see show multicast next-hops on page 391.

show multicast next-hops detail

```
user@host> show multicast next-hops detail
Family: INET
ID          Refcount KRefcount Downstream interface
1048577            2         1 fe-0/1/2.0-(1048574)
                               ge-0/2/3.0-(1048576)
```

show multicast pim-to-igmp-proxy

List of Syntax
Syntax on page 393
Syntax (EX Series Switch and the QFX Series) on page 393

Syntax
show multicast pim-to-igmp-proxy
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show multicast pim-to-igmp-proxy
<instance *instance-name*>

Release Information
Command introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 9.6 for EX Series switches.
instance option introduced in Junos OS Release 10.3.
instance option introduced in Junos OS Release 10.3 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.

Options
none—Display configuration information about PIM-to-IGMP message translation for all routing instances.

instance *instance-name*—(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- *Configuring PIM-to-IGMP and PIM-to-MLD Message Translation*

List of Sample Output
show multicast pim-to-igmp-proxy on page 394
show multicast pim-to-igmp-proxy instance on page 394

Output Fields
Table 29 on page 393 describes the output fields for the **show multicast pim-to-igmp-proxy** command. Output fields are listed in the order in which they appear.

Table 29: show multicast pim-to-igmp-proxy Output Fields

Field Name	Field Description
Instance	Routing instance. Default instance is **master** (inet.0 routing table).
Proxy state	State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: **enabled** or **disabled**.

Table 29: show multicast pim-to-igmp-proxy Output Fields *(continued)*

Field Name	Field Description
interface-name	Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured.

Sample Output

show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

show multicast pim-to-igmp-proxy instance

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```

show multicast pim-to-mld-proxy

List of Syntax
Syntax on page 395
Syntax (EX Series Switch and the QFX Series) on page 395

Syntax
show multicast pim-to-mld-proxy
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show multicast pim-to-mld-proxy
<instance *instance-name*>

Release Information
Command introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 9.6 for EX Series switches.
instance option introduced in Junos OS Release 10.3.
instance option introduced in Junos OS Release 10.3 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.

Options
none—Display configuration information about PIM-to-MLD message translation for all routing instances.

instance *instance-name*—(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast pim-to-mld-proxy on page 396
show multicast pim-to-mld-proxy instance on page 396

Output Fields
Table 30 on page 395 describes the output fields for the **show multicast pim-to-mld-proxy** command. Output fields are listed in the order in which they appear.

Table 30: show multicast pim-to-mld-proxy Output Fields

Field Name	Field Description
Proxy state	State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: **enabled** or **disabled**.
interface-name	Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured.

Sample Output

show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

show multicast pim-to-mld-proxy instance

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```

show multicast route

List of Syntax

Syntax on page 397
Syntax (EX Series Switch and the QFX Series) on page 397

Syntax

show multicast route
<brief | detail | extensive | summary>
<active | all | inactive>
<group *group*>
<inet | inet6>
<instance *instance name*>
<logical-system (all | *logical-system-name*)>
<*regular-expression*>
<source-prefix *source-prefix*>

Syntax (EX Series Switch and the QFX Series)

show multicast route
<brief | detail | extensive | summary>
<active | all | inactive>
<group *group*>
<inet | inet6>
<instance *instance name*>
<*regular-expression*>
<source-prefix *source-prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.

Description

Display the entries in the IP multicast forwarding table. You can display similar information with the **show route table inet.1** command.

Options

none—Display standard information about all entries in the multicast forwarding table for all routing instances.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

active | all | inactive—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.

group *group*—(Optional) Display the cache entries for a particular group.

inet | inet6—(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display entries in the multicast forwarding table for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

regular-expression—(Optional) Display information about the multicast forwarding table entries that match a UNIX OS-style regular expression.

source-prefix *source-prefix*—(Optional) Display the cache entries for a particular source prefix.

Required Privilege Level view

List of Sample Output

show multicast route on page 399
show multicast route (Bidirectional PIM) on page 400
show multicast route brief on page 400
show multicast route detail on page 400
show multicast route extensive (Bidirectional PIM) on page 401
show multicast route instance <instance-name> on page 402
show multicast route summary on page 402

Output Fields Table 31 on page 398 describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

Table 31: show multicast route Output Fields

Field Name	Field Description	Level of Output
family	IPv4 address family (**INET**) or IPv6 address family (**INET6**).	All levels
Group	Group address. For any-source multicast routes, for example for bidirectional PIM, the group address includes the prefix length.	All levels
Source	Prefix and length of the source as it is in the multicast forwarding table.	All levels
Incoming interface list	List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	All levels
Upstream interface	Name of the interface on which the packet with this source prefix is expected to arrive.	All levels
Downstream interface list	List of interface names to which the packet with this source prefix is forwarded.	All levels
Number of outgoing interfaces	Total number of outgoing interfaces for each (S,G) entry.	**extensive**
Session description	Name of the multicast session.	**detail extensive**
Statistics	Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobits per second packet forwarding statistic queries fails or times out, the statistics field displays **Forwarding statistics are not available**. NOTE: On QFX Series switches, this field does not report valid statistics.	**detail extensive**

Table 31: show multicast route Output Fields *(continued)*

Field Name	Field Description	Level of Output
Next-hop ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the **show multicast nexthops** command.	**detail extensive**
Incoming interface list ID	For bidirectional PIM, incoming interface list identifier. Identifiers for interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	**detail extensive**
Upstream protocol	The protocol that maintains the active multicast forwarding route for this group or source. When the **show multicast route extensive** command is used with the **display-origin-protocol** option, the field name is only **Protocol** and not **Upstream Protocol**. However, this field also displays the protocol that installed the active route.	**detail extensive**
Route type	Type of multicast route. Values can be (S,G) or (*,G).	**summary**
Route state	Whether the group is **Active** or **Inactive**.	**summary extensive**
Route count	Number of multicast routes.	**summary**
Forwarding state	Whether the prefix is pruned or forwarding.	**extensive**
Cache lifetime/timeout	Number of seconds until the prefix is removed from the multicast forwarding table. A value of **never** indicates a permanent forwarding entry. A value of **forever** indicates routes that do not have keepalive times.	**extensive**
Wrong incoming interface notifications	Number of times that the upstream interface was not available.	**extensive**
Uptime	Time since the creation of a multicast route.	**extensive**

Sample Output

show multicast route

```
user@host> show multicast route
Family: INET

Group: 228.0.0.0
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0

Group: 239.1.1.1
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
```

```
        so-1/0/0.0

Group: 239.1.1.1
    Source: 10.255.70.15/32
    Upstream interface: so-1/0/0.0
    Downstream interface list:
        mt-1/1/0.1081344

Family: INET6
```

show multicast route (Bidirectional PIM)

```
user@host> show multicast route
Family: INET

Group: 224.1.1.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0

Group: 224.1.3.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0

Group: 225.1.1.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0

Group: 225.1.3.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0
Family: INET6
```

show multicast route brief

The output for the **show multicast route brief** command is identical to that for the **show multicast route** command. For sample output, see show multicast route on page 399 or show multicast route (Bidirectional PIM) on page 400.

show multicast route detail

```
user@host> show multicast route detail
Family: INET

Group: 228.0.0.0
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Unknown
```



```
    Statistics: 8 kBps, 100 pps, 45272 packets
    Next-hop ID: 262142
    Upstream protocol: PIM

Group: 239.1.1.1
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Administratively Scoped
    Statistics: 0 kBps, 0 pps, 13404 packets
    Next-hop ID: 262142
    Upstream protocol: PIM

Group: 239.1.1.1
    Source: 10.255.70.15/32
    Upstream interface: so-1/0/0.0
    Downstream interface list:
        mt-1/1/0.1081344
    Session description: Administratively Scoped
    Statistics: 46 kBps, 1000 pps, 921077 packets

    Next-hop ID: 262143
    Upstream protocol: PIM

Family: INET6
```

show multicast route extensive (Bidirectional PIM)

```
user@host> show multicast route extensive
Family: INET

Group: 224.1.1.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0
    Number of outgoing interfaces: 1
    Session description: NOB Cross media facilities
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097153
    Incoming interface list ID: 585
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Group: 224.1.3.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0
    Number of outgoing interfaces: 1
    Session description: NOB Cross media facilities
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097153
    Incoming interface list ID: 589
    Upstream protocol: PIM
```

```
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Family: INET6
```

show multicast route instance <instance-name>

```
user@host> show multicast route instance v1 extensive
Instance: v1 Family: INET

Group: 224.1.1.1
    Source: (null)/0
    Upstream interface: fe-1/3/0.111
    Downstream interface list:
        lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
    Number of outgoing interfaces: 3

Group: 224.1.1.2
    Source: (null)/0
    Upstream interface: fe-1/3/0.111
    Downstream interface list:
        lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
    Number of outgoing interfaces: 3

Group: 224.1.1.3
    Source: (null)/0
    Upstream interface: fe-1/3/0.111
    Downstream interface list:
        lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
    Number of outgoing interfaces: 3

Instance: v1 Family: INET6
```

show multicast route summary

```
user@host>show multicast route summary
Instance: master Family: INET

Route type   Route state    Route count
(S,G)             Active             2
(S,G)             Inactive           3

Instance: master Family: INET6
```

show multicast rpf

List of Syntax
Syntax on page 403
Syntax (EX Series Switch and the QFX Series) on page 403

Syntax

show multicast rpf
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*prefix*>
<summary>

Syntax (EX Series Switch and the QFX Series)

show multicast rpf
<inet | inet6>
<instance *instance-name*>
<*prefix*>
<summary>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display information about multicast reverse-path-forwarding (RPF) calculations.

Options

none—Display RPF calculation information for all supported address families.

inet | inet6—(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about multicast RPF calculations for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix—(Optional) Display the RPF calculation information for the specified prefix.

summary—(Optional) Display a summary of all multicast RPF information.

Required Privilege Level

view

List of Sample Output

show multicast rpf on page 404
show multicast rpf inet6 on page 405
show multicast rpf prefix on page 406
show multicast rpf summary on page 406

Output Fields Table 32 on page 404 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

Table 32: show multicast rpf Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Source prefix	Prefix and length of the source as it exists in the multicast forwarding table.
Protocol	How the route was learned.
Interface	Upstream RPF interface. NOTE: The displayed interface information does not apply to bidirectional PIM RP addresses. This is because the **show multicast rpf** command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF interface information, always use the **show pim join extensive** command when bidirectional PIM is configured.
Neighbor	Upstream RPF neighbor. NOTE: The displayed neighbor information does not apply to bidirectional PIM. This is because the **show multicast rpf** command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF neighbor information, always use the **show pim join extensive** command when bidirectional PIM is configured.

Sample Output

show multicast rpf

```
user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
    Protocol: Static

10.255.14.132/32
    Protocol: Direct
    Interface: lo0.0

10.255.245.91/32
    Protocol: IS-IS
    Interface: so-1/1/1.0
    Neighbor: 192.168.195.21

127.0.0.1/32
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0
```

```
Neighbor: 192.168.14.254

192.168.0.0/16
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

192.168.14.0/24
Protocol: Direct
Interface: fxp0.0

192.168.14.132/32
Protocol: Local

192.168.195.20/30
Protocol: Direct
Interface: so-1/1/1.0

192.168.195.22/32
Protocol: Local

192.168.195.36/30
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 192.168.195.21
```

show multicast rpf inet6

```
user@host> show multicast rpf inet6

Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128
    Protocol: Direct
    Interface: lo0.0

::10.255.245.91/128
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.20/126
Protocol: Direct
Interface: so-1/1/1.0

::192.168.195.22/128
Protocol: Local

::192.168.195.36/126
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.76/126
Protocol: Direct
Interface: fe-2/2/0.0

::192.168.195.77/128
Protocol: Local
```

```
fe80::/64
Protocol: Direct
Interface: so-1/1/1.0

fe80::290:69ff:fe0c:993a/128
Protocol: Local

fe80::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0

ff02::2/128
Protocol: PIM

ff02::d/128
Protocol: PIM
```

show multicast rpf prefix

```
user@host> show multicast rpf ff02::/16

Multicast RPF table: inet6.0, 13 entries

ff02::2/128
    Protocol: PIM

ff02::d/128
    Protocol: PIM

...
```

show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```

show multicast scope

List of Syntax
Syntax on page 407
Syntax (EX Series Switch and the QFX Series) on page 407

Syntax
show multicast scope
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show multicast scope
<inet | inet6>
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display administratively scoped IP multicast information.

Options
none—Display standard information about administratively scoped multicast information for all supported address families in all routing instances.

inet | inet6—(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display administratively scoped information for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast scope on page 408
show multicast scope inet on page 408
show multicast scope inet6 on page 408

Output Fields
Table 33 on page 407 describes the output fields for the **show multicast scope** command. Output fields are listed in the approximate order in which they appear.

Table 33: show multicast scope Output Fields

Field Name	Field Description
Scope name	Name of the multicast scope.
Group Prefix	Range of multicast groups that are scoped.
Interface	Interface that is the boundary of the administrative scope.

Table 33: show multicast scope Output Fields *(continued)*

Field Name	Field Description
Resolve Rejects	Number of kernel resolve rejects.

Sample Output

show multicast scope

```
user@host> show multicast scope
                                                          Resolve
Scope name           Group Prefix       Interface         Rejects
232-net              232.232.0.0/16     fe-0/0/0.1              0
local                239.255.0.0/16     fe-0/0/0.1              0
local                ff05::/16          fe-0/0/0.1              0
larry                ff05::1234/128     fe-0/0/0.1              0
```

show multicast scope inet

```
user@host> show multicast scope inet
                                                          Resolve
Scope name           Group Prefix       Interface         Rejects
232-net              232.232.0.0/16     fe-0/0/0.1              0
local                239.255.0.0/16     fe-0/0/0.1              0
```

show multicast scope inet6

```
user@host> show multicast scope inet6
                                                          Resolve
Scope name           Group Prefix       Interface         Rejects
local                ff05::/16          fe-0/0/0.1              0
larry                ff05::1234/128     fe-0/0/0.1              0
```

show multicast sessions

List of Syntax
Syntax on page 409
Syntax (EX Series Switch and the QFX Series) on page 409

Syntax
show multicast sessions
<brief | detail | extensive>
<logical-system (all | *logical-system-name*)>
<*regular-expression*>

Syntax (EX Series Switch and the QFX Series)
show multicast sessions
<brief | detail | extensive>
<*regular-expression*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display information about announced IP multicast sessions.

Options
none—Display standard information about all multicast sessions for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

regular-expression—(Optional) Display information about announced sessions that match a UNIX-style regular expression.

Required Privilege Level
view

List of Sample Output
show multicast sessions on page 410
show multicast sessions regular-expression detail on page 410

Output Fields
Table 34 on page 409 describes the output fields for the **show multicast sessions** command. Output fields are listed in the approximate order in which they appear.

Table 34: show multicast sessions Output Fields

Field Name	Field Description
session-name	Name of the known announced multicast sessions.

Sample Output

show multicast sessions

```
user@host> show multicast sessions
1-Department of Biological Sciences, LSU
  ...
  Monterey Bay - DockCam
  Monterey Bay - JettyCam
  Monterey Bay - StandCam
  Monterey DockCam
  Monterey DockCam / ROV cam
  ...
  NASA TV (MPEG-1)
  ...
  UO Broadcast - NASA Videos - 25 Years of Progress
  UO Broadcast - NASA Videos - Journey through the Solar System
  UO Broadcast - NASA Videos - Life in the Universe
  UO Broadcast - NASA Videos - Nasa and the Airplane
  UO Broadcasts OPB's Oregon Story
  UO DOD News Clips
  UO Medical Management of Biological Casualties (1)
  UO Medical Management of Biological Casualties (2)
  UO Medical Management of Biological Casualties (3)
...
  376 active sessions.
```

show multicast sessions regular-expression detail

```
user@host> show multicast sessions "NASA TV" detail
SDP Version: 0  Originated by: -@128.223.83.33
 Session: NASA TV (MPEG-1)
 Description: NASA television in MPEG-1 format, provided by Private University.
Please contact the UO if you have problems with this feed.
 Email: Your Name Here <multicast@lists.private.edu>
 Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 224.2.231.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 128.223.91.191 live
Attribute: fmtp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106
 Connection Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
Attribute: rtpmap:103 L16/22050/2
```

```
Attribute: rtpmap:104 L16/22050

  1 matching sessions.
```

show multicast usage

List of Syntax
Syntax on page 412
Syntax (EX Series Switch and the QFX Series) on page 412

Syntax

show multicast usage
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show multicast usage
<brief | detail>
<inet | inet6>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups.

Options

none—Display multicast usage information for all supported address families for all routing instances.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show multicast usage on page 413
show multicast usage brief on page 413
show multicast usage instance on page 413
show multicast usage detail on page 414

Output Fields

Table 35 on page 413 describes the output fields for the **show multicast usage** command. Output fields are listed in the approximate order in which they appear.

Table 35: show multicast usage Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Group	Group address.
Sources	Number of sources.
Packets	Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays **unavailable**.
Bytes	Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays **unavailable**.
Prefix	IP address.
/len	Prefix length.
Groups	Number of multicast groups.

Sample Output

show multicast usage

```
user@host> show multicast usage
Group           Sources Packets             Bytes
228.0.0.0       1       52847               4439148
239.1.1.1       2       13450               1125530


Prefix          /len Groups Packets             Bytes
10.255.14.144   /32  2      66254               5561304
10.255.70.15    /32  1      43                  3374...
```

show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see show multicast usage on page 413.

show multicast usage instance

```
user@host> show multicast usage instance VPN-A
Group           Sources Packets             Bytes
224.2.127.254   1       5538                509496
224.0.1.39      1       13                  624
224.0.1.40      1       13                  624

Prefix          /len Groups Packets             Bytes
192.168.195.34  /32  1      5538                509496
10.255.14.30    /32  1      13                  624
```

```
10.255.245.91   /32  1      13                  624
...
```

show multicast usage detail

```
user@host> show multicast usage detail
Group           Sources Packets              Bytes
228.0.0.0       1       53159                4465356
    Source: 10.255.14.144   /32 Packets: 53159 Bytes: 4465356
239.1.1.1       2       13450                1125530
    Source: 10.255.14.144   /32 Packets: 13407 Bytes: 1122156
    Source: 10.255.70.15    /32 Packets: 43 Bytes: 3374


Prefix          /len Groups Packets              Bytes
10.255.14.144   /32  2      66566               5587512
    Group: 228.0.0.0       Packets: 53159 Bytes: 4465356
    Group: 239.1.1.1       Packets: 13407 Bytes: 1122156
10.255.70.15    /32  1      43                  3374
    Group: 239.1.1.1       Packets: 43 Bytes: 3374
```

show pim bootstrap

List of Syntax
Syntax on page 415
Syntax (EX Series Switch and the QFX Series) on page 415

Syntax
show pim bootstrap
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show pim bootstrap
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
instance option introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers.

Options
none—Display PIM bootstrap router information for all routing instances.

instance *instance-name*—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show pim bootstrap on page 416
show pim bootstrap instance on page 416

Output Fields
Table 36 on page 415 describes the output fields for the **show pim bootstrap** command. Output fields are listed in the approximate order in which they appear.

Table 36: show pim bootstrap Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
BSR	Bootstrap router.
Pri	Priority of the routing device as elected to be the bootstrap router.
Local address	Local routing device address.
Pri	Local routing device address priority to be elected as the bootstrap router.

Table 36: show pim bootstrap Output Fields *(continued)*

Field Name	Field Description
State	Local routing device election state: **Candidate**, **Elected**, or **Ineligible**.
Timeout	How long until the local routing device declares the bootstrap router to be unreachable, in seconds.

Sample Output

show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master

BSR                    Pri Local address           Pri State      Timeout
None                     0 10.255.71.46              0 InEligible       0
feco:1:1:1:1:0:aff:785c 34 feco:1:1:1:1:0:aff:7c12   0 InEligible       0
```

show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A

BSR             Pri Local address   Pri State      Timeout
None              0 192.168.196.105   0 InEligible       0
```

show pim interfaces

List of Syntax
Syntax on page 417
Syntax (EX Series Switch and the QFX Series) on page 417

Syntax
show pim interfaces
<inet | inet6>
<instance (*instance-name* | all)>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show pim interfaces
<inet | inet6>
<instance (*instance-name* | all)>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Commmand introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.
Support for the **instance all** option added in Junos OS Release 12.1.

Description
Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured.

Options
none—Display interface information for all family addresses for the main instance.

inet | inet6—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.

instance (*instance-name* | all)—(Optional) Display information about interfaces for a specific PIM-enabled routing instance or for all routing instances.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show pim interfaces on page 418

Output Fields
Table 37 on page 417 describes the output fields for the **show pim interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 37: show pim interfaces Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Name	Interface name.
State	State of the interface. The state also is displayed in the **show interfaces** command.

Table 37: show pim interfaces Output Fields *(continued)*

Field Name	Field Description
Mode	PIM mode running on the interface: • **B**—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers. • **S**—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic. • **Dense**—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.) • **Sparse-Dense**—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as **dense** is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as **sparse** is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules. (Not supported on QFX Series.) When sparse-dense mode is configured, the output includes both **S** and **D**. When bidirectional-sparse mode is configured, the output includes **S** and **B**. When bidirectional-sparse-dense mode is configured, the output includes **B**, **S**, and **D**.
IP	Version number of the address family on the interface: **4** (IPv4) or **6** (IPv6).
V	PIM version running on the interface: **1** or **2**.
State	State of PIM on the interface: • **Active**—Bidirectional mode is enabled on the interface and on all PIM neighbors. • **DR**—Designated router. • **NotCap**—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol. • **NotDR**—Not the designated router. • **P2P**—Point to point.
NbrCnt	Number of neighbors that have been seen on the interface.
JoinCnt(sg)	Number of (s,g) join messages that have been seen on the interface.
JointCnt(*g)	Number of (*,g) join messages that have been seen on the interface.
DR address	Address of the designated router.

Sample Output

show pim interfaces

```
user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
```

```
Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name               Stat Mode IP V State        NbrCnt JoinCnt(sg/*g) DR address
ge-0/3/0.0         Up   S     4 2 NotDR,NotCap      1 0/0            40.0.0.3
ge-0/3/3.50        Up   S     4 2 DR,NotCap         1 9901/100       50.0.0.2
ge-0/3/3.51        Up   S     4 2 DR,NotCap         1 0/0            51.0.0.2
pe-1/2/0.32769     Up   S     4 2 P2P,NotCap        0 0/0
```

show pim join

List of Syntax
Syntax on page 420
Syntax (EX Series Switch and the QFX Series) on page 420

Syntax

show pim join
<brief | detail | extensive | summary>
<bidirectional | dense | sparse>
<exact>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*range*>
<rp *ip-address/prefix* | source *ip-address/prefix*>
<sg | star-g>

Syntax (EX Series Switch and the QFX Series)

show pim join
<brief | detail | extensive | summary>
<dense | sparse>
<exact>
<inet | inet6>
<instance *instance-name*>
<*range*>
<rp *ip-address/prefix* | source *ip-address/prefix*>
<sg | star-g>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
summary option introduced in Junos OS Release 9.6.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Support for bidirectional PIM added in Junos OS Release 12.1.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Multiple new filter options introduced in Junos OS Release 13.2.

Description

Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.

For bidirectional PIM, display information about PIM group ranges (*,G-range) for each active bidirectional RP group range, in addition to each of the joined (*,G) routes.

Options

none—Display the standard information about PIM groups for all supported family addresses for all routing instances.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

bidirectional | dense | sparse—(Optional) Display information about PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.

exact—(Optional) Display information about only the group that exactly matches the specified group address.

inet | inet6—(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about groups for the specified PIM-enabled routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

range—(Optional) Address range of the group, specified as ***prefix/prefix-length***.

rp *ip-address/prefix* | source *ip-address/prefix*—(Optional) Display information about the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

sg | star-g—(Optional) Display information about PIM (S,G) or (*,G) entries.

Required Privilege Level view

Related Documentation
- clear pim join on page 333

List of Sample Output
show pim join summary on page 425
show pim join (PIM Sparse Mode) on page 425
show pim join (Bidirectional PIM) on page 425
show pim join inet6 on page 426
show pim join inet6 star-g on page 426
show pim join instance <instance-name> on page 426
show pim join detail on page 427
show pim join extensive (PIM Sparse Mode) on page 427
show pim join extensive (Bidirectional PIM) on page 428
show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP) on page 429
show pim join instance <instance-name> extensive on page 430
show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 430
show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 431

Output Fields Table 38 on page 421 describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

Table 38: show pim join Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	**brief detail extensive summary** none
Family	Name of the address family: **inet** (IPv4) or **inet6** (IPv6).	**brief detail extensive summary** none

Table 38: show pim join Output Fields *(continued)*

Field Name	Field Description	Level of Output
Route type	Type of multicast route: (S,G) or (*,G).	**summary**
Route count	Number of (S,G) routes and number of (*,G) routes.	**summary**
R	Rendezvous Point Tree.	**brief detail extensive** none
S	Sparse.	**brief detail extensive** none
W	Wildcard.	**brief detail extensive** none
Group	Group address.	**brief detail extensive** none
Bidirectional group prefix length	For bidirectional PIM, length of the IP prefix for RP group ranges.	All levels
Source	Multicast source: • ***** (wildcard value) • ***ipv4-address*** • ***ipv6-address***	**brief detail extensive** none
RP	Rendezvous point for the PIM group.	**brief detail extensive** none
Flags	PIM flags: • **bidirectional**—Bidirectional mode entry. • **dense**—Dense mode entry. • **rptree**—Entry is on the rendezvous point tree. • **sparse**—Sparse mode entry. • **spt**—Entry is on the shortest-path tree for the source. • **wildcard**—Entry is on the shared tree.	**brief detail extensive** none
Upstream interface	RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*,G). For bidirectional PIM, **RP Link** means that the interface is directly connected to a subnet that contains a phantom RP address. A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.	**brief detail extensive** none

Table 38: show pim join Output Fields *(continued)*

Field Name	Field Description	Level of Output
Upstream neighbor	Information about the upstream neighbor: **Direct**, **Local**, **Unknown**, or a specific IP address. For bidirectional PIM, **Direct** means that the interface is directly connected to a subnet that contains a phantom RP address. The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.	**extensive**
Upstream state	Information about the upstream interface: • **Join to RP**—Sending a join to the rendezvous point. • **Join to Source**—Sending a join to the source. • **Local RP**—Sending neither join messages nor prune messages toward the RP, because this routing device is the rendezvous point. • **Local Source**—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device. • **Prune to RP**—Sending a prune to the rendezvous point. • **Prune to Source**—Sending a prune to the source. NOTE: RP group range entries have **None** in the **Upstream state** field because RP group ranges do not trigger actual PIM join messages between routing devices.	**extensive**

Table 38: show pim join Output Fields *(continued)*

Field Name	Field Description	Level of Output
Downstream neighbors	Information about downstream interfaces: • **Interface**—Interface name for the downstream neighbor. A pseudo PIM-SM interface appears for all IGMP-only interfaces. A pseudo multipoint LDP (M-LDP) interface appears on ingress root nodes in M-LDP point-to-multipoint LSPs with inband signaling. • **Interface address**—Address of the downstream neighbor. • **State**—Information about the downstream neighbor: **join** or **prune**. • **Flags**—PIM join flags: **R (RPtree)**, **S (Sparse)**, **W (Wildcard)**, or **zero**. • **Uptime**—Time since the downstream interface joined the group. • **Time since last Join**—Time since the last join message was received from the downstream interface. • **Time since last Prune**—Time since the last prune message was received from the downstream interface.	**extensive**
Number of downstream interfaces	Total number of outgoing interfaces for each (S,G) entry.	**extensive**
Assert Timeout	Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null.	**extensive**
Keepalive timeout	Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, **Keepalive timeout** is **Infinity**.	**extensive**
Uptime	Time since the creation of (S,G) or (*,G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*,G) state.	**extensive**
Bidirectional accepting interfaces	Interfaces on the routing device that forward bidirectional PIM traffic. The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (**DF Winner**), or the interface is the reverse path forwarding (RPF) interface toward the RP (**RPF**).	**extensive**

Sample Output

show pim join summary

```
user@host> show pim join summary
Instance: PIM.master Family: INET

Route type            Route count
(s,g)                 2
(*,g)                 1

Instance: PIM.master Family: INET6
```

show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
    Bidirectional group prefix length: 24
    Source: *
```

```
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join inet6

```
user@host> show pim join inet6
Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

 Group: ff04::e000:101
     Source: *
     RP: ::46.0.0.13
     Flags: sparse,rptree,wildcard
     Upstream interface: Local

 Group: ff04::e000:101
     Source: ::1.1.1.1
     Flags: sparse
     Upstream interface: unknown (no neighbor)

 Group: ff04::e800:101
     Source: ::1.1.1.1
     Flags: sparse
     Upstream interface: unknown (no neighbor)

 Group: ff04::e800:101
     Source: ::1.1.1.2
     Flags: sparse
     Upstream interface: unknown (no neighbor)
```

show pim join inet6 star-g

```
user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

 Group: ff04::e000:101
     Source: *
     RP: ::46.0.0.13
     Flags: sparse,rptree,wildcard
     Upstream interface: Local
```

show pim join instance <instance-name>

```
user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
    Source: *
    RP: 10.10.47.100
```

```
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 235.1.1.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0

Group: 235.1.1.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join detail

```
user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: so-1/0/0.0
            10.111.10.2 State: Join Flags: SRW Timeout: 174
            Uptime: 00:03:49 Time since last Join: 00:01:49
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: SRW  Timeout: Infinity
```

```
            Uptime: 00:03:49 Time since last Join: 00:01:49
        Number of downstream interfaces: 2

    Group: 239.1.1.1
        Source: 10.255.14.144
        Flags: sparse,spt
        Upstream interface: Local
        Upstream neighbor: Local
        Upstream state: Local Source, Local RP
        Keepalive timeout: 344
        Uptime: 00:03:49
        Downstream neighbors:
            Interface: so-1/0/0.0
                10.111.10.2 State: Join Flags: S Timeout: 174
                Uptime: 00:03:49 Time since last Prune: 00:01:49
            Interface: mt-1/1/0.32768
                10.10.47.100 State: Join Flags: S   Timeout: Infinity
                Uptime: 00:03:49 Time since last Prune: 00:01:49
        Number of downstream interfaces: 2

    Group: 239.1.1.1
        Source: 10.255.70.15
        Flags: sparse,spt
        Upstream interface: so-1/0/0.0
        Upstream neighbor: 10.111.10.2
        Upstream state: Local RP, Join to Source
        Keepalive timeout: 344
        Uptime: 00:03:49
        Downstream neighbors:
            Interface: Pseudo-GMP
                fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
            Interface: so-1/0/0.0 (pruned)
                10.111.10.2 State: Prune Flags: SR Timeout: 174
                Uptime: 00:03:49 Time since last Prune: 00:01:49
            Interface: mt-1/1/0.32768
                10.10.47.100 State: Join Flags: S   Timeout: Infinity
                Uptime: 00:03:49 Time since last Prune: 00:01:49
        Number of downstream interfaces: 3

    Instance: PIM.master Family: INET6
    R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM)

```
    user@host> show pim join extensive
    Instance: PIM.master Family: INET
    R = Rendezvous Point Tree, S = Sparse, W = Wildcard

    Group: 224.1.1.0
        Bidirectional group prefix length: 24
        Source: *
        RP: 10.10.13.2
        Flags: bidirectional,rptree,wildcard
        Upstream interface: ge-0/0/1.0
        Upstream neighbor: 10.10.1.2
        Upstream state: None
        Uptime: 00:03:49
        Bidirectional accepting interfaces:
            Interface: ge-0/0/1.0    (RPF)
            Interface: lo0.0         (DF Winner)
        Number of downstream interfaces: 0
```

```
Group: 225.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.10.1.2
    Upstream state: None
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
     Downstream neighbors:
         Interface: lt-1/0/10.24
             10.0.24.4 State: Join   RW  Timeout: 185
         Interface: lt-1/0/10.23
             10.0.23.3 State: Join   RW  Timeout: 184
     Number of downstream interfaces: 2

Group: 225.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)
    Upstream neighbor: Direct
    Upstream state: Local RP
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
        Interface: xe-4/1/0.0    (DF Winner)
    Number of downstream interfaces: 0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)
    Upstream neighbor: Direct
    Upstream state: Local RP
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
        Interface: xe-4/1/0.0    (DF Winner)
    Number of downstream interfaces: 0
```

show pim join instance <instance-name> extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
    Source: *
    RP: 10.10.47.100
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: mt-1/1/0.32768
            10.10.47.101 State: Join Flags: SRW Timeout: 156
            Uptime: 00:03:49 Time since last Join: 00:01:49
    Number of downstream interfaces: 1

Group: 235.1.1.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0
    Upstream neighbor: 10.111.30.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156
    Uptime: 00:14:52

Group: 235.1.1.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0
    Upstream neighbor: 10.111.20.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156
    Uptime: 00:14:52
```

show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:55
    Downstream neighbors:
        Interface: Pseudo-MLDP
        Interface: lt-1/2/0.25
            1.2.5.2 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
    Source: 192.168.219.11
    Flags: sparse,spt
```

```
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:41
    Downstream neighbors:
        Interface: Pseudo-MLDP

Group: 232.1.1.3
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:41
    Downstream neighbors:
        Interface: Pseudo-MLDP

Group: 232.2.2.2
    Source: 1.2.7.7
    Flags: sparse,spt
    Upstream interface: lt-1/2/0.27
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:25
    Downstream neighbors:
        Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
    Source: abcd::1:2:7:7
    Flags: sparse,spt
    Upstream interface: lt-1/2/0.27
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:26
    Downstream neighbors:
        Interface: Pseudo-MLDP
```

show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 227.1.1.1
    Source: *
    RP: 1.1.1.1
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 11:31:33
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: SRW  Timeout: Infinity
```

```
            Uptime: 11:31:33 Time since last Join: 11:31:32

Group: 232.1.1.1
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <1.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.2
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <1.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30
    Downstream neighbors:
        Interface: lt-1/2/0.14
            1.1.4.4 State: Join Flags: S Timeout: 177
            Uptime: 11:30:33 Time since last Join: 00:00:33
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.3
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <1.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.2.2.2
    Source: 1.2.7.7
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
```

```
    Upstream neighbor: MLDP LSP root <1.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:30
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
    Source: abcd::1:2:7:7
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <1.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: fe-1/3/0.0
            fe80::21f:12ff:fea5:c4db State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32
```

show pim neighbors

List of Syntax
Syntax on page 434
Syntax (EX Series Switch and the QFX Series) on page 434

Syntax
show pim neighbors
<brief | detail>
<inet | inet6>
<instance (*instance-name* | all)>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show pim neighbors
<brief | detail>
<inet | inet6>
<instance (*instance-name* | all)>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.
Support for the **instance all** option added in Junos OS Release 12.1.

Description
Display information about Protocol Independent Multicast (PIM) neighbors.

Options
none—(Same as **brief**) Display standard information about PIM neighbors for all supported family addresses for the main instance.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.

instance (*instance-name* | all)—(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show pim neighbors on page 436
show pim neighbors brief on page 436
show pim neighbors instance on page 436
show pim neighbors detail on page 436
show pim neighbors detail (With BFD) on page 437

Output Fields
Table 39 on page 435 describes the output fields for the **show pim neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 39: show pim neighbors Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Interface	Interface through which the neighbor is reachable.	All levels
Neighbor addr	Address of the neighboring PIM routing device.	All levels
IP	IP version: **4** or **6**.	All levels
V	PIM version running on the neighbor: **1** or **2**.	All levels
Mode	PIM mode of the neighbor: **Sparse**, **Dense**, **SparseDense**, or **Unknown**. When the neighbor is running PIM version 2, this mode is always **Unknown**.	All levels
Option	Can be one or more of the following: • **B**—Bidirectional Capable. • **H**—Hello Option Holdtime. • **G**—Generation Identifier. • **P**—Hello Option DR Priority. • **L**—Hello Option LAN Prune Delay.	**brief** none
Uptime	Time the neighbor has been operational since the PIM process was last initialized, in the format ***dd*:*hh*:*mm*:*ss* ago** for less than a week and ***nwnd*:*hh*:*mm*:*ss* ago** for more than a week.	All levels
Address	Address of the neighboring PIM routing device.	**detail**
BFD	Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: **Enabled**, **Operational state is up**, or **Disabled**.	**detail**
Hello Option Holdtime	Time for which the neighbor is available, in seconds. The range of values is **0** through **65,535.**	**detail**
Hello Default Holdtime	Default holdtime and the time remaining if the **holdtime** option is not in the received hello message.	**detail**
Hello Option DR Priority	Designated router election priority. The range of values is **0** through **255**.	**detail**
Hello Option Generation ID	9-digit or 10-digit number used to tag hello messages.	**detail**
Hello Option Bi-Directional PIM supported	Neighbor can process bidirectional PIM messages.	**detail**
Hello Option LAN Prune Delay	Time to wait before the neighbor receives prune messages, in the format **delay *nnn* ms override *nnnn* ms**.	**detail**

Table 39: show pim neighbors Output Fields *(continued)*

Field Name	Field Description	Level of Output
Join Suppression supported	Neighbor is capable of join suppression.	**detail**
Rx Join	Information about joins received from the neighbor. • **Group**—Group addresses in the join message. • **Source**—Address of the source in the join message. • **Timeout**—Time for which the join is valid.	**detail**

Sample Output

show pim neighbors

```
user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface           IP V Mode        Option      Uptime Neighbor addr
so-1/0/0.0           4 2             HPLG      00:07:10 10.111.10.2
```

show pim neighbors brief

The output for the **show pim neighbors brief** command is identical to that for the **show pim neighbors** command. For sample output, see show pim neighbors on page 436.

show pim neighbors instance

```
user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface           IP V Mode        Option      Uptime Neighbor addr
at-0/3/1.0           4 2             HPLG      00:07:54 10.111.30.2
mt-1/1/0.32768       4 2             HPLG      00:07:22 10.10.47.101
so-1/0/1.0           4 2             HPLG      00:07:50 10.111.20.2
```

show pim neighbors detail

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: ge-0/0/1.0

    Address: 10.10.1.1, IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, tsg
Join Count: 2
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 2053759302
        Hello Option Bi-Directional PIM supported
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                      Join Suppression supported
```

```
    Address: 10.10.1.2, IPv4, PIM v2, sg Join Count: 0, tsg Join Count: 2
        BFD: Disabled
        Hello Option Holdtime: 105 seconds 93 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1734018161
        Hello Option Bi-Directional PIM supported
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                      Join Suppression supported

Interface: lo0.0

    Address: 10.255.179.246,    IPv4, PIM v2, Mode: SparseDense, sg Join Count:
0, tsg Join Count: 0
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1997462267
        Hello Option Bi-Directional PIM supported
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                      Join Suppression supported
```

show pim neighbors detail (With BFD)

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: fe-1/0/0.0
    Address: 192.168.11.1,        IPv4, PIM v2, Mode: Sparse
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 836607909
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

    Address: 192.168.11.2,        IPv4, PIM v2
        BFD: Enabled, Operational state is up
        Hello Default Holdtime: 105 seconds 104 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1907549685
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Interface: fe-1/0/1.0
    Address: 192.168.12.1,        IPv4, PIM v2
        BFD: Disabled
        Hello Default Holdtime: 105 seconds 80 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1971554705
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

show pim rps

List of Syntax
Syntax on page 438
Syntax (EX Series Switch and the QFX Series) on page 438

Syntax

show pim rps
<brief | detail | extensive>
<*group-address*>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show pim rps
<brief | detail | extensive>
<*group-address*>
<inet | inet6>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.

Description

Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs).

Options

none—Display standard information about PIM RPs for all groups and family addresses for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

group-address—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.

inet | inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about RPs for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *Example: Configuring Bidirectional PIM*

List of Sample Output

show pim rps on page 441
show pim rps brief on page 441
show pim rps <group-address> (Bidirectional PIM) on page 441

show pim rps <group-address> (PIM Dense Mode) on page 441
show pim rps <group-address> (SSM Range Without asm-override-ssm Configured) on page 441
show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP) on page 442
show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP) on page 442
show pim rps instance on page 442
show pim rps extensive (PIM Sparse Mode) on page 442
show pim rps extensive (Bidirectional PIM) on page 443
show pim rps extensive (PIM Anycast RP in Use) on page 443

Output Fields Table 40 on page 439 describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

Table 40: show pim rps Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Family or Address family	Name of the address family: **inet** (IPv4) or **inet6** (IPv6).	All levels
RP address	Address of the rendezvous point.	All levels
Type	Type of RP: • **auto-rp**—Address of the RP known through the Auto-RP protocol. • **bootstrap**—Address of the RP known through the bootstrap router protocol (BSR). • **embedded**—Address of the RP known through an embedded RP (IPv6). • **static**—Address of RP known through static configuration.	**brief** none
Holdtime	How long to keep the RP active, with time remaining, in seconds.	All levels
Timeout	How long until the local routing device determines the RP to be unreachable, in seconds.	All levels
Groups	Number of groups currently using this RP.	All levels
Group prefixes	Addresses of groups that this RP can span.	**brief** none
Learned via	Address and method by which the RP was learned.	**detail extensive**
Mode	The PIM mode of the RP: bidirectional or sparse. If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats.	All levels
Time Active	How long the RP has been active, in the format ***hh:mm:ss***.	**detail extensive**

Table 40: show pim rps Output Fields *(continued)*

Field Name	Field Description	Level of Output
Device Index	Index value of the order in which Junos OS finds and initializes the interface. For bidirectional RPs, the **Device Index** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Subunit	Logical unit number of the interface. For bidirectional RPs, the **Subunit** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Interface	Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively. For bidirectional RPs, the **Interface** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Group Ranges	Addresses of groups that this RP spans.	**detail extensive** *group-address*
Active groups using RP	Number of groups currently using this RP.	**detail extensive**
total	Total number of active groups for this RP.	**detail extensive**
Register State for RP	Current register state for each group: • **Group**—Multicast group address. • **Source**—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively: • **First Hop**—PIM-designated routing device that sent the Register message (the source address in the IP header). • **RP Address**—RP to which the Register message was sent (the destination address in the IP header). • **State**: On the designated router: ◦ **Send**—Sending Register messages. ◦ **Probe**—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages. ◦ **Suppress**—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to **Probe** state. • On the RP: ◦ **Receive**—Receiving Register messages.	**extensive**
Anycast-PIM rpset	If anycast RP is configured, the addresses of the RPs in the set.	**extensive**
Anycast-PIM local address used	If anycast RP is configured, the local address used by the RP.	**extensive**

Table 40: show pim rps Output Fields *(continued)*

Field Name	Field Description	Level of Output

Anycast-PIM Register State	If anycast RP is configured, the current register state for each group:
	• **Group**—Multicast group address.
	• **Source**—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively.
	• **Origin**—How the information was obtained:
	  • **DIRECT**—From a local attachment
	  • **MSDP**—From the Multicast Source Discovery Protocol (MSDP)
	  • **DR**—From the designated router
		extensive

RP selected	For sparse mode and bidirectional mode, the identity of the RP for the specified group address.	*group-address*

Sample Output

show pim rps

```
user@host> show pim rps
Instance: PIM.master
Address family INET
RP address      Type        Mode   Holdtime Timeout Groups Group prefixes
10.10.1.3       static      bidir       150    None      2 224.1.3.0/24
                                                           225.1.3.0/24
10.10.13.2      static      bidir       150    None      2 224.1.1.0/24
                                                           225.1.1.0/24
```

show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see show pim rps on page 441.

show pim rps <group-address> (Bidirectional PIM)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

224.1.0.0/16
        11.4.12.75 (Bidirectional)

RP selected: 11.4.12.75
```

show pim rps <group-address> (PIM Dense Mode)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Dense Mode active for group 224.1.1.1
```

show pim rps <group-address> (SSM Range Without asm-override-ssm Configured)

```
user@host> show pim rps 224.1.1.1
```

```
Instance: PIM.master

Source-specific Mode (SSM) active for group 224.1.1.1
```

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16
        11.4.12.75

RP selected: 11.4.12.75
```

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16
        11.4.12.75 (Bidirectional)

RP selected: (null)
```

show pim rps instance

```
user@host> show pim rps instance VPN-A
Instance: PIM.VPN-A
Address family INET
RP address              Type        Holdtime Timeout Groups Group prefixes
10.10.47.100            static             0    None      1 224.0.0.0/4

Address family INET6
```

show pim rps extensive (PIM Sparse Mode)

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.255.245.91
Learned via: static configuration
Time Active: 00:05:48
Holdtime: 45 with 36 remaining
Device Index: 122
Subunit: 32768
Interface: pd-6/0/0.32768
Group Ranges:
    224.0.0.0/4, 36s remaining
Active groups using RP:
    225.1.1.1

    total 1 groups active

Register State for RP:
Group           Source          FirstHop        RP Address      State    Timeout
225.1.1.1       192.168.195.78  10.255.14.132   10.255.245.91   Receive        0
```

show pim rps extensive (Bidirectional PIM)

```
user@host> show pim rps extensive
Instance: PIM.master
Address family INET

RP: 10.10.1.3
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
        224.1.3.0/24
        225.1.3.0/24

RP: 10.10.13.2
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
        224.1.1.0/24
        225.1.1.0/24
```

show pim rps extensive (PIM Anycast RP in Use)

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.10.10.2
Learned via: static configuration
Time Active: 00:54:52
Holdtime: 0
Device Index: 130
Subunit: 32769
Interface: pimd.32769
Group Ranges:
    224.0.0.0/4
Active groups using RP:
    224.10.10.10

    total 1 groups active

Anycast-PIM rpset:
           10.100.111.34
           10.100.111.17
           10.100.111.55

Anycast-PIM local address used: 10.100.111.1
Anycast-PIM Register State:
Group                   Source                                  Origin
224.1.1.1               10.10.95.2                              DIRECT
224.1.1.2               10.10.95.2                              DIRECT
224.10.10.10            10.10.70.1                              MSDP
224.10.10.11            10.10.70.1                              MSDP
224.20.20.1             10.10.71.1                              DR

Address family INET6

Anycast-PIM rpset:
```

```
           ab::1
           ab::2
Anycast-PIM local address used: cd::1

Anycast-PIM Register State:
Group                   Source                                Origin
::224.1.1.1             ::10.10.95.2                          DIRECT
::224.1.1.2             ::10.10.95.2                          DIRECT
::224.20.20.1           ::10.10.71.1                          DR
```

show pim source

List of Syntax
Syntax on page 445
Syntax (EX Series Switch and the QFX Series) on page 445

Syntax

show pim source
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*source-prefix*>

Syntax (EX Series Switch and the QFX Series)

show pim source
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<*source-prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state.

Options

none—Display standard information about the PIM RPF state for all supported family addresses for all routing instances.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

source-prefix—(Optional) Display the state for source RPF states in the given range.

Required Privilege Level

view

List of Sample Output

show pim source on page 446
show pim source brief on page 446
show pim source detail on page 446
show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 447

Output Fields

Table 41 on page 446 describes the output fields for the **show pim source** command. Output fields are listed in the approximate order in which they appear.

Table 41: show pim source Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Source	Address of the source or reverse path.
Prefix/length	Prefix and prefix length for the route used to reach the RPF address.
Upstream Protocol	
Upstream interface	RPF interface toward the source address. A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.
Upstream Neighbor	Address of the RPF neighbor used to reach the source address. The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.

Sample Output

show pim source

```
user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
    Prefix 10.255.14.144/32
    Upstream interface Local
    Upstream neighbor Local

Source 10.255.70.15
    Prefix 10.255.70.15/32
    Upstream interface so-1/0/0.0
    Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6
```

show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see show pim source on page 446.

show pim source detail

```
user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
    Prefix 10.255.14.144/32
    Upstream interface Local
    Upstream neighbor Local
    Active groups:228.0.0.0
        239.1.1.1
```

```
        239.1.1.1

Source 10.255.70.15
    Prefix 10.255.70.15/32
    Upstream interface so-1/0/0.0
    Upstream neighbor 10.111.10.2
    Active groups:239.1.1.1

Instance: PIM.master Family: INET6
```

show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim source
Instance: PIM.master Family: INET

Source 1.1.1.1
    Prefix 1.1.1.1/32
    Upstream interface Local
    Upstream neighbor Local

Source 1.2.7.7
    Prefix 1.2.7.0/24
    Upstream protocol MLDP
    Upstream interface Pseudo MLDP
    Upstream neighbor MLDP LSP root <1.1.1.2>

Source 192.168.219.11
    Prefix 192.168.219.0/28
    Upstream protocol MLDP
    Upstream interface Pseudo MLDP
    Upstream neighbor MLDP LSP root <1.1.1.2>

Instance: PIM.master Family: INET6

Source abcd::1:2:7:7
    Prefix abcd::1:2:7:0/120
    Upstream protocol MLDP
    Upstream interface Pseudo MLDP
    Upstream neighbor MLDP LSP root <1.1.1.2>
```

show pim statistics

List of Syntax
Syntax on page 448
Syntax (EX Series Switch and the QFX Series) on page 448

Syntax

```
show pim statistics
<inet |inet6)>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show pim statistics
<inet |inet6)>
<instance instance-name>
<interface interface-name>
```

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.

Description
Display Protocol Independent Multicast (PIM) statistics.

Options
none—Display PIM statistics.

inet | inet6—(Optional) Display IPv4 or IPv6 PIM statistics, respectively.

instance *instance-name*—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM).

interface *interface-name*—(Optional) Display statistics about the specified interface.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear pim statistics on page 337

List of Sample Output
show pim statistics on page 455
show pim statistics inet interface <interface-name> on page 457
show pim statistics inet6 interface <interface-name> on page 457
show pim statistics instance <instance-name> on page 458
show pim statistics interface <interface-name> on page 459

Output Fields
Table 42 on page 449 describes the output fields for the **show pim statistics** command. Output fields are listed in the approximate order in which they appear.

Table 42: show pim statistics Output Fields

Field Name	Field Description
Instance	Name of the routing instance. This field only appears if you specify an interface, for example: • **inet interface *interface-name*** • **inet6 interface *interface-name*** • **interface *interface-name***
Family	Output is for IPv4 or IPv6 PIM statistics. **INET** indicates IPv4 statistics, and **INET6** indicates IPv6 statistics. This field only appears if you specify an interface, for example: • **inet interface *interface-name*** • **inet6 interface *interface-name*** • **interface *interface-name***
PIM statistics	PIM statistics for all interfaces or for the specified interface.
PIM message type	Message type for which statistics are displayed.
Received	Number of received statistics.
Sent	Number of messages sent of a certain type.
Rx errors	Number of received packets that contained errors.
V2 Hello	PIM version 2 hello packets.
V2 Register	PIM version 2 register packets.
V2 Register Stop	PIM version 2 register stop packets.
V2 Join Prune	PIM version 2 join and prune packets.
V2 Bootstrap	PIM version 2 bootstrap packets.
V2 Assert	PIM version 2 assert packets.
V2 Graft	PIM version 2 graft packets.
V2 Graft Ack	PIM version 2 graft acknowledgment packets.
V2 Candidate RP	PIM version 2 candidate RP packets.

Table 42: show pim statistics Output Fields *(continued)*

Field Name	Field Description
V2 State Refresh	PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh. State refresh is an extension to PIM-DM. It not supported in Junos OS.
V2 DF Election	PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election.
V1 Query	PIM version 1 query packets.
V1 Register	PIM version 1 register packets.
V1 Register Stop	PIM version 1 register stop packets.
V1 Join Prune	PIM version 1 join and prune packets.
V1 RP Reachability	PIM version 1 RP reachability packets.
V1 Assert	PIM version 1 assert packets.
V1 Graft	PIM version 1 graft packets.
V1 Graft Ack	PIM version 1 graft acknowledgment packets.
AutoRP Announce	Auto-RP announce packets.
AutoRP Mapping	Auto-RP mapping packets.
AutoRP Unknown type	Auto-RP packets with an unknown type.
Anycast Register	Auto-RP announce packets.
Anycast Register Stop	Auto-RP announce packets.
Global Statistics	Summary of PIM statistics for all interfaces.
Hello dropped on neighbor policy	Number of hello packets dropped because of a configured neighbor policy.
Unknown type	Number of PIM control packets received with an unknown type.
V1 Unknown type	Number of PIM version 1 control packets received with an unknown type.
Unknown Version	Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.

Table 42: show pim statistics Output Fields *(continued)*

Field Name	Field Description
Neighbor unknown	Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.
Bad Length	Number of PIM control packets received for which the packet size does not match the PIM length field in the packet.
Bad Checksum	Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet.
Bad Receive If	Number of PIM control packets received on an interface that does not have PIM configured.
Rx Bad Data	Number of PIM control packets received that contain data for TCP Bad register packets.
Rx Intf disabled	Number of PIM control packets received on an interface that has PIM disabled.
Rx V1 Require V2	Number of PIM version 1 control packets received on an interface configured for PIM version 2.
Rx V2 Require V1	Number of PIM version 2 control packets received on an interface configured for PIM version 1.
Rx Register not RP	Number of PIM register packets received when the routing device is not the RP for the group.
Rx Register no route	Number of PIM register packets received when the RP does not have a unicast route back to the source.
Rx Register no decap if	Number of PIM register packets received when the RP does not have a de-encapsulation interface.
Null Register Timeout	Number of NULL register timeout packets.
RP Filtered Source	Number of PIM packets received when the routing device has a source address filter configured for the RP.
Rx Unknown Reg Stop	Number of register stop messages received with an unknown type.
Rx Join/Prune no state	Number of join and prune messages received for which the routing device has no state.
Rx Join/Prune on upstream if	Number of join and prune messages received on the interface used to reach the upstream routing device, toward the RP.
Rx Join/Prune for invalid group	Number of join or prune messages received for invalid multicast group addresses.

Table 42: show pim statistics Output Fields *(continued)*

Field Name	Field Description
Rx Join/Prune messages dropped	Number of join and prune messages received and dropped.
Rx sparse join for dense group	Number of PIM sparse mode join messages received for a group that is configured for dense mode.
Rx Graft/Graft Ack no state	Number of graft and graft acknowledgment messages received for which the router or switch has no state.
Rx Graft on upstream if	Number of graft messages received on the interface used to reach the upstream routing device, toward the RP.
Rx CRP not BSR	Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap.
Rx BSR when BSR	Number of BSR messages received in which the PIM message type is Bootstrap.
Rx BSR not RPF if	Number of BSR messages received on an interface that is not the RPF interface.
Rx unknown hello opt	Number of PIM hello packets received with options that Junos OS does not support.
Rx data no state	Number of PIM control packets received for which the routing device has no state for the data type.
Rx RP no state	Number of PIM control packets received for which the routing device has no state for the RP.
Rx aggregate	Number of PIM aggregate MDT packets received.
Rx malformed packet	Number of PIM control packets received with a malformed IP unicast or multicast address family.
No RP	Number of PIM control packets received with no RP address.
No register encap if	Number of PIM register packets received when the first-hop routing device does not have an encapsulation interface.
No route upstream	Number of PIM control packets received when the routing device does not have a unicast route to the the interface used to reach the upstream routing device, toward the RP.
Nexthop Unusable	Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down.
RP mismatch	Number of PIM control packets received for which the routing device has an RP mismatch.

Table 42: show pim statistics Output Fields *(continued)*

Field Name	Field Description
RP mode mismatch	RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages.
RPF neighbor unknown	Number of PIM control packets received for which the routing device has an unknown RPF neighbor for the source.
Rx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Tx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Embedded-RP invalid addr	Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.
Embedded-RP limit exceed	Number of times the limit configured with the **maximum-rps** statement is exceeded. The **maximum-rps** statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.
Embedded-RP added	Number of packets in which the embedded RP for IPv6 is added. The following receive events trigger extraction of an IPv6 embedded RP address on the routing device: • Multicast Listener Discovery (MLD) report for an embedded RP multicast group address • PIM join message with an embedded RP multicast group address • Static embedded RP multicast group address associated with an interface • Packets sent to an embedded RP multicast group address received on the DR An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.
Embedded-RP removed	Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.
Rx Register msgs filtering drop	Number of received register messages dropped because of a filter configured for PIM register messages.
Tx Register msgs filtering drop	Number of register messages dropped because of a filter configured for PIM register messages.
Rx Bidir Join/Prune on non-Bidir if	Error counter for join and prune messages received on non-bidirectional PIM interfaces.

Table 42: show pim statistics Output Fields *(continued)*

Field Name	Field Description
Rx Bidir Join/Prune on non-DF if	Error counter for join and prune messages received on non-designated forwarder interfaces.
V4 (S,G) Maximum	Maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted.
V4 (S,G) Accepted	Number of accepted (S,G) IPv4 multicast routes.
V4 (S,G) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device).
V4 (S,G) Log Interval	Time (in seconds) between consecutive log messages.
V6 (S,G) Maximum	Maximum number of (S,G) IPv6 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted.
V6 (S,G) Accepted	Number of accepted (S,G) IPv6 multicast routes.
V6 (S,G) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv6 multicast routes accepted by the device).
V6 (S,G) Log Interval	Time (in seconds) between consecutive log messages.
V4 (grp-prefix, RP) Maximum	Maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.
V4 (grp-prefix, RP) Accepted	Number of accepted group-to-RP IPv4 multicast mappings.
V4 (grp-prefix, RP) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device).
V4 (grp-prefix, RP) Log Interval	Time (in seconds) between consecutive log messages.
V6 (grp-prefix, RP) Maximum	Maximum number of group-to RP IPv6 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.
V6 (grp-prefix, RP) Accepted	Number of accepted group-to-RP IPv6 multicast mappings.

Table 42: show pim statistics Output Fields *(continued)*

Field Name	Field Description
V6 (grp-prefix, RP) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv6 multicast mappings accepted by the device).
V6 (grp-prefix, RP) Log Interval	Time (in seconds) between consecutive log messages.
V4 Register Maximum	Maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP.
V4 Register Accepted	Number of accepted IPv4 PIM registers.
V4 Register Threshold	Threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device).
V4 Register Log Interval	Time (in seconds) between consecutive log messages.
V6 Register Maximum	Maximum number of IPv6 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP.
V6 Register Accepted	Number of accepted IPv6 PIM registers.
V6 Register Threshold	Threshold at which a warning message is logged (percentage of the maximum number of IPv6 PIM registers accepted by the device).
V6 Register Log Interval	Time (in seconds) between consecutive log messages.

Sample Output

show pim statistics

```
user@host> show pim statistics
PIM Message type        Received       Sent  Rx errors
V2 Hello                      15         32          0
V2 Register                    0        362          0
V2 Register Stop             483          0          0
V2 Join Prune                 18        518          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V2 State Refresh               0          0          0
V2 DF Election                 0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
```

```
V1 Register Stop              0          0          0
V1 Join Prune                 0          0          0
V1 RP Reachability            0          0          0
V1 Assert                     0          0          0
V1 Graft                      0          0          0
V1 Graft Ack                  0          0          0
AutoRP Announce               0          0          0
AutoRP Mapping                0          0          0
AutoRP Unknown type           0
Anycast Register              0          0          0
Anycast Register Stop         0          0          0

Global Statistics

Hello dropped on neighbor policy              0
Unknown type                                  0
V1 Unknown type                               0
Unknown Version                               0
ipv4 BSR pkt drop due to excessive rate       0
ipv6 BSR pkt drop due to excessive rate       0
Neighbor unknown                              0
Bad Length                                    0
Bad Checksum                                  0
Bad Receive If                                0
Rx Bad Data                                   0
Rx Intf disabled                              0
Rx V1 Require V2                              0
Rx V2 Require V1                              0
Rx Register not RP                            0
Rx Register no route                          0
Rx Register no decap if                       0
Null Register Timeout                         0
RP Filtered Source                            0
Rx Unknown Reg Stop                           0
Rx Join/Prune no state                        0
Rx Join/Prune on upstream if                  0
Rx Join/Prune for invalid group               5
Rx Join/Prune messages dropped                0
Rx sparse join for dense group                0
Rx Graft/Graft Ack no state                   0
Rx Graft on upstream if                       0
Rx CRP not BSR                                0
Rx BSR when BSR                               0
Rx BSR not RPF if                             0
Rx unknown hello opt                          0
Rx data no state                              0
Rx RP no state                                0
Rx aggregate                                  0
Rx malformed packet                           0
Rx illegal TTL                                0
Rx illegal destination address                0
No RP                                         0
No register encap if                          0
No route upstream                             0
Nexthop Unusable                              0
RP mismatch                                   0
RP mode mismatch                              0
RPF neighbor unknown                          0
Rx Joins/Prunes filtered                      0
Tx Joins/Prunes filtered                      0
Embedded-RP invalid addr                      0
```

```
Embedded-RP limit exceed                    0
Embedded-RP added                           0
Embedded-RP removed                         0
Rx Register msgs filtering drop             0
Tx Register msgs filtering drop             0
Rx Bidir Join/Prune on non-Bidir if         0
Rx Bidir Join/Prune on non-DF if            0
```

Sample Output

show pim statistics inet interface <interface-name>

```
user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          4          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
AutoRP Announce                0          0          0
AutoRP Mapping                 0          0          0
AutoRP Unknown type            0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0
```

Sample Output

show pim statistics inet6 interface <interface-name>

```
user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          4          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
```

```
Anycast Register              0          0          0
Anycast Register Stop         0          0          0
```

show pim statistics instance <instance-name>

```
user@host> show pim statistics instance VPN-A
PIM Message type        Received       Sent  Rx errors
V2 Hello                      31         37          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0         16          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V2 State Refresh               0          0          0
V2 DF Election                 0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
AutoRP Announce                0          0          0
AutoRP Mapping                 0          0          0
AutoRP Unknown type            0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0

Global Statistics

Hello dropped on neighbor policy              0
Unknown type                                  0
V1 Unknown type                               0
Unknown Version                               0
Neighbor unknown                              0
Bad Length                                    0
Bad Checksum                                  0
Bad Receive If                                0
Rx Bad Data                                   0
Rx Intf disabled                              0
Rx V1 Require V2                              0
Rx V2 Require V1                              0
Rx Register not RP                            0
Rx Register no route                          0
Rx Register no decap if                       0
Null Register Timeout                         0
RP Filtered Source                            0
Rx Unknown Reg Stop                           0
Rx Join/Prune no state                        0
Rx Join/Prune on upstream if                  0
Rx Join/Prune for invalid group               0
Rx Join/Prune messages dropped                0
Rx sparse join for dense group                0
Rx Graft/Graft Ack no state                   0
Rx Graft on upstream if                       0
Rx CRP not BSR                                0
Rx BSR when BSR                               0
```

```
Rx BSR not RPF if                          0
Rx unknown hello opt                       0
Rx data no state                           0
Rx RP no state                             0
Rx aggregate                               0
Rx malformed packet                        0
Rx illegal TTL                             0
Rx illegal destination address             0
No RP                                      0
No register encap if                       0
No route upstream                         28
Nexthop Unusable                           0
RP mismatch                                0
RP mode mismatch                           0
RPF neighbor unknown                       0
Rx Joins/Prunes filtered                   0
Tx Joins/Prunes filtered                   0
Embedded-RP invalid addr                   0
Embedded-RP limit exceed                   0
Embedded-RP added                          0
Embedded-RP removed                        0
Rx Register msgs filtering drop            0
Tx Register msgs filtering drop            0
Rx Bidir Join/Prune on non-Bidir if        0
Rx Bidir Join/Prune on non-DF if           0
V4 (S,G) Maximum                          10
V4 (S,G) Accepted                          9
V4 (S,G) Threshold                        80
V4 (S,G) Log Interval                     80
V6 (S,G) Maximum                           8
V6 (S,G) Accepted                          8
V6 (S,G) Threshold                        50
V6 (S,G) Log Interval                    100
V4 (grp-prefix, RP) Maximum              100
V4 (grp-prefix, RP) Accepted               5
V4 (grp-prefix, RP) Threshold             80
V4 (grp-prefix, RP) Log Interval          10
V6 (grp-prefix, RP) Maximum               20
V6 (grp-prefix, RP) Accepted               0
V6 (grp-prefix, RP) Threshold             90
V6 (grp-prefix, RP) Log Interval          20
V4 Register Maximum                      100
V4 Register Accepted                      10
V4 Register Threshold                     80
V4 Register Log Interval                  10
V6 Register Maximum                       20
V6 Register Accepted                       0
V6 Register Threshold                     90
V6 Register Log Interval                  20
```

Sample Output

show pim statistics interface <interface-name>

```
user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          3          0
```

```
V2 Register                   0          0          0
V2 Register Stop              0          0          0
V2 Join Prune                 0          0          0
V2 Bootstrap                  0          0          0
V2 Assert                     0          0          0
V2 Graft                      0          0          0
V2 Graft Ack                  0          0          0
V2 Candidate RP               0          0          0
V1 Query                      0          0          0
V1 Register                   0          0          0
V1 Register Stop              0          0          0
V1 Join Prune                 0          0          0
V1 RP Reachability            0          0          0
V1 Assert                     0          0          0
V1 Graft                      0          0          0
V1 Graft Ack                  0          0          0
AutoRP Announce               0          0          0
AutoRP Mapping                0          0          0
AutoRP Unknown type           0
Anycast Register              0          0          0
Anycast Register Stop         0          0          0

Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          3          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0
```

show system statistics igmp

List of Syntax

Syntax

```
show system statistics igmp
```

Syntax (EX Series Switches)

```
show system statistics igmp
<all-members>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics igmp
<all-chassis | all-lcc | lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics igmp
<all-chassis | all-lcc | lcc number | sfc number>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description

Display system-wide Internet Group Management Protocol (IGMP) statistics.

Options

none—Display system statistics for IGMP.

all-chassis—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system statistics for IGMP for all the routers in the chassis.

all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for IGMP for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for all connected T1600 or T4000 LCCs.

all-members—(EX4200 switches only) (Optional) Display IGMP statistics for all members of the Virtual Chassis configuration.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for IGMP for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches only) (Optional) Display IGMP statistics for the local Virtual Chassis member.

member *member-id*—(EX4200 switches only) (Optional) Display IGMP statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

scc—(TX Matrix routers only) (Optional) Display system statistics for IGMP for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display system statistics for IGMP for the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system statistics igmp** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation
- *Routing Matrix with a TX Matrix Plus Router Solutions Page*

List of Sample Output
show system statistics igmp on page 462
show system statistics igmp (EX Series Switches) on page 462
show system statistics igmp (TX Matrix Plus Router) on page 463

Sample Output

show system statistics igmp

```
user@host> show system statistics igmp
igmp:
        17178 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent
```

show system statistics igmp (EX Series Switches)

```
user@host> show system statistics igmp
```

```
igmp:
         0 messages received
         0 messages received with too few bytes
         0 messages received with bad checksum
         0 membership queries received
         0 membership queries received with invalid fields
         0 membership reports received
         0 membership reports received with invalid fields
         0 membership reports received for groups to which we belong
         0 Membership reports sent
```

show system statistics igmp (TX Matrix Plus Router)

```
user@host> show system statistics igmp
sfc0-re0:
--------------------------------------------------------------------------
igmp:
        0 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent

lcc0-re0:
--------------------------------------------------------------------------
igmp:
        0 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent

lcc1-re0:
--------------------------------------------------------------------------
igmp:
        0 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent

lcc2-re0:
--------------------------------------------------------------------------
igmp:
        0 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
```

```
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent

lcc3-re0:
--------------------------------------------------------------------------
igmp:
        0 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent
```

test msdp

Syntax test msdp (dependent-peers *prefix* | rpf-peer *originator*)
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description Find Multicast Source Discovery Protocol (MSDP) peers.

Options

dependent-peers *prefix*—Find downstream dependent MSDP peers.

rpf-peer *originator*—Find the MSDP reverse-path-forwarding (RPF) peer for the originator.

instance *instance-name*—(Optional) Find MDSP peers for the specified routing instance.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

List of Sample Output test msdp dependent-peers on page 465

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

test msdp dependent-peers

```
user@host> test msdp dependent-peers 10.0.0.1/24
```

PART 4

Index

- Index on page 469

Index

Symbols

A

B

C

D

E

F

G

K

L

M

N

O

P

S

T

V

W