JunosE™ Software for E Series™ Broadband Services Routers

Multicast Routing Configuration Guide

Release

14.1.x

Published: 2012-12-12

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

JunosE™ Software for E Series™ Broadband Services Routers Multicast Routing Configuration Guide
Release 14.1.x

Revision History
December 2012—FRS JunosE 14.1.x

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.

Abbreviated Table of Contents

Table of Contents

List of Figures

List of Tables

About the Documentation

E Series and JunosE Documentation and Release Notes

For a list of related JunosE documentation, see http://www.juniper.net/techpubs/software/index.html.

If the information in the latest release notes differs from the information in the documentation, follow the *JunosE 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/.

Audience

This guide is intended for experienced system and network specialists working with Juniper Networks E Series Broadband Services Routers in an Internet access environment.

E Series and JunosE Text and Syntax Conventions

Table 1 on page xx defines notice icons used in this documentation.

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.

Table 2 on page xx defines text and syntax conventions that we use throughout the E Series and JunosE documentation.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents commands and keywords in text.	• Issue the **clock source** command. • Specify the keyword **exp-msg**.
Bold text like this	Represents text that the user must type.	**host1(config)#traffic class low-loss1**

`Fixed-width text like this`	Represents information as displayed on your terminal's screen.	**host1#show ip ospf 2**
		`Routing Process OSPF 2 with Router ID 5.5.0.250`
		`Router is an Area Border Router (ABR)`

Italic text like this	• Emphasizes words. • Identifies variables. • Identifies chapter, appendix, and book names.	• There are two levels of access: *user* and *privileged*. • *clusterId, ipAddress.* • *Appendix A, System Specifications*
Plus sign (+) linking key names	Indicates that you must press two or more keys simultaneously.	Press Ctrl + b.
Syntax Conventions in the Command Reference Guide		
Plain text like this	Represents keywords.	terminal length
Italic text like this	Represents variables.	*mask*, *accessListName*

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

Convention	Description	Examples
| (pipe symbol)	Represents a choice to select one keyword or variable to the left or to the right of this symbol. (The keyword or variable can be either optional or required.)	diagnostic | line
[ ] (brackets)	Represent optional keywords or variables.	[internal | external]
[ ]* (brackets and asterisk)	Represent optional keywords or variables that can be entered more than once.	[level1 | level2 | l1]*
{ } (braces)	Represent required keywords or variables.	{ permit | deny } { in | out } { *clusterId* | *ipAddress* }

Obtaining Documentation

To obtain the most current version of all Juniper Networks technical documentation, see the Technical Documentation page on the Juniper Networks Web site at http://www.juniper.net/.

To download complete sets of technical documentation to create your own documentation CD-ROMs or DVD-ROMs, see the Portable Libraries page at

http://www.juniper.net/techpubs/resources/index.html

Copies of the Management Information Bases (MIBs) for a particular software release are available for download in the software image bundle from the Juniper Networks Web site athttp://www.juniper.net/.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation to better meet your needs. Send your comments to **techpubs-comments@juniper.net**, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version

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: https://www.juniper.net/alerts/
- 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

Internet Protocol Version 4

CHAPTER 1

Configuring IPv4 Multicast

IPv4 multicast enables a device to send packets to a group of hosts rather than to a list of individual hosts. This chapter describes how to configure IP multicast on the E Series router; it contains the following sections:

IPv4 Multicast Overview

IPv4 defines three types of addresses: *unicast*, *broadcast*, and *multicast*. Each type of address enables a device to send datagrams to selected recipients:

- A unicast address enables a device to send a datagram to a single recipient.
- A broadcast address enables a device to send a datagram to all hosts on a subnetwork.
- A multicast address enables a device to send a datagram to a specified set of hosts, known as a multicast group, in different subnetworks.

Multicast IP packets contain a class D address in the Destination Address fields of their headers. A class D address is the IP address of a multicast group.

IP multicast improves network efficiency by enabling a host to transmit a datagram to a targeted group of receivers. For example, for a host to send a large video clip to a group of selected recipients would be time-consuming to unicast the datagram to each recipient individually. If the host broadcasts the video clip throughout the network, network resources are not available for other tasks. The host uses only the resources it needs when multicasting the datagram.

Routers use multicast routing algorithms to determine the best route and transmit multicast datagrams throughout the network. E Series routers support a number of IP multicast protocols on virtual routers (VRs). Each VR handles the interoperability of IP multicast protocols automatically. To start multicast operation on a VR, you access the context for that VR and configure the desired protocols on the selected interfaces. Table 3 on page 4 describes the function of each protocol that the router supports.

Table 3: Function of Multicast Protocols on a Router

Protocol	Function
Internet Group Membership Protocol (IGMP)	Discovers hosts that belong to multicast group.
Protocol Independent Multicast Protocol (PIM)	Discovers other multicast routers to receive multicast packets.
Distance Vector Multicast Routing Protocol (DVMRP)	Routes multicast datagrams within autonomous systems.
BGP Multicasting Protocol	Routes multicast datagrams between autonomous systems.

The router supports up to 16,384 multicast forwarding entries (multicast routes) at any time.

This topic discusses the following sections:

- Reverse-Path Forwarding on page 5
- Multicast Packet Forwarding on page 5

Reverse-Path Forwarding

IP multicasting uses reverse path forwarding (RPF) to verify that a router receives a multicast packet on the correct incoming interface. The RPF algorithm enables a router to accept a multicast datagram only on the interface from which the router sends a unicast datagram to the source of the multicast datagram.

When the router receives a multicast datagram from a source for a group, the router verifies that the packet was received on the correct RPF interface. If the packet was not received on the correct interface, the router discards the packet. Only packets received on the correct RPF interface are considered for forwarding to downstream receivers.

When operating in sparse-mode, the routers perform an RPF lookup to identify the upstream router from which to request the data and then send join messages for the multicast stream only to that router.

When operating in dense-mode, routers that have multiple paths to the source of the multicast stream initially receive the same stream on more than one interface. In this case, the routers perform an RPF lookup to identify multicast data streams that are not arriving on the best path and send prune messages to terminate these flows.

The RPF lookup need not always be towards the source of the multicast stream. The lookup is done towards the source only when the router is using a source-rooted tree to receive the multicast stream. If the router uses a shared tree instead, the RPF lookup is toward a rendezvous point and not toward the source of the multicast stream.

Multicast Packet Forwarding

Multicast packet forwarding is based on the source (S) of the multicast packet and the destination multicast group address (G). For each (S,G) pair, the router accepts multicast packets on an incoming interface (IIF), which satisfies the RPF check (RPF-IIF). The router drops packets received on IIFs other than the RPF-IIF and notifies the routing protocols that a packet was received on the wrong interface.

The router forwards packets received on the RPF-IIF to a list of outgoing interfaces (OIFs). The list of OIFs is determined by the exchange of routing information and local group membership information. The router maintains mappings of (S,G, IIF) to {OIF1, OIF2...} in the multicast routing table.

You can enable two or more multicast protocols on an IIF. However, only one protocol can forward packets on that IIF. The protocol that forwards packets on an IIF *owns* that IIF. A multicast protocol that owns an IIF also owns the (S,G) entry in the multicast routing table.

Related Documentation

- Configuring IPv4 Multicast Attributes on page 7
- IPv4 Multicast Platform Considerations on page 6
- IPv4 Multicast References on page 6

IPv4 Multicast Platform Considerations

For information about modules that support IP multicasting on the ERX7xx models, ERX14xx models, and the Juniper Networks ERX310 Broadband Services Router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support IP multicasting.

For information about modules that support IP multicasting on the Juniper Networks E120 and E320 Broadband Services Routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support IP multicasting.

Related Documentation

- IPv4 Multicast Overview on page 4
- IPv4 Multicast References on page 6

IPv4 Multicast References

For more information about IP multicast, see the following resources:

- A " traceroute" Facility for IP Multicast—draft-ietf-idmr-traceroute-ipm-07.txt (January 2001 expiration)
- RFC 2858—Multiprotocol Extensions for BGP-4 (June 2000)
- RFC 2932—IPv4 Multicast Routing MIB (October 2000)
- RFC 3292—General Switch Management Protocol (GSMP) V3 (June 2002)

NOTE: **IETF drafts are valid for only 6 months from the date of issuance. They must be considered as works in progress. Refer to the IETF Web site at http://www.ietf.org for the latest drafts.**

Related Documentation

- IPv4 Multicast Overview on page 4
- IPv4 Multicast Platform Considerations on page 6

Switch Fabric Bandwidth Configuration

By default, the switch fabric for the ERX1440, ERX310, E120, and E320 Broadband Services Routers uses a bandwidth weighting ratio of 15:2 for multicast-to-unicast weighted round robin (WRR). In the absence of strict-priority traffic, and when both unicast and multicast traffic compete for switch fabric bandwidth, the switch fabric allocates 15/17ths of the available bandwidth to multicast traffic and 2/17ths of the available bandwidth to unicast traffic.

You can use the **fabric weights** command to change the ratio for multicast to unicast traffic on the router switch fabric. For more information about the **fabric weights** command, see *Configuring the Switch Fabric Bandwidth* in the *JunosE System Basics Configuration Guide*.

Related Documentation

- IPv4 Multicast Overview on page 4
- IPv6 Multicast Overview on page 161

- fabric weights

Configuring IPv4 Multicast Attributes

You can configure IPv4 multicast attributes with the following tasks:

Before You Begin Configuring IPv4 Multicast

You can configure multicasting on IPv4 and IPv6 interfaces.

For information about configuring IP and IPv6 interfaces, see *JunosE IP, IPv6, and IGP Configuration Guide*.

Enabling IP Multicast

In this implementation, IP multicast works on virtual routers (VRs). By default, IP multicast is disabled on a VR. To enable IP multicast on a VR, access the context for a VR, and then issue the **ip multicast-routing** command.

```
host1(config)#ip multicast-routing
```

You can use the **no** version to disable IP multicast routing on the VR (the default).

Defining Static Routes for Reverse-Path Forwarding

Use the **ip rpf-route** command to define reverse-path forwarding (RPF) to verify that a router receives a multicast packet on the correct incoming interface.

```
host1(config)#ip rpf-route 11.1.0.0 255.255.0.0 atm4/1.1 56 tag 25093
```

You can use the **no** version to disable IP multicast routing on the VR (the default). In the disabled state, all multicast protocols are disabled, and the VR forwards no multicast packets.

Enabling and Disabling RPF Checks

By default, the router accepts multicast packets for each Source, Group (S,G) pair on an incoming interface (IIF), which satisfies the RPF check (RPF-IIF). When the router performs RPF checks, only the interface that first accepts traffic for an (S,G) pair accepts subsequent traffic for that pair. If traffic stops arriving on that interface and starts arriving on another interface, the router does not accept or forward the traffic.

Some network configurations require the router to accept traffic on any interface. To do so, you can disable the RPF check on a specified set of (S,G) pairs by issuing the **ip multicast-routing disable-rpf-check** command.

```
host1(config)#ip multicast-routing disable-rpf-check boston-list
```

You can use the **no** version to restore the default, in which the router performs RPF checks for all (S,G) pairs.

When you disable RPF checks, the router accepts multicast packets for (S,G) pairs on any incoming interface. When the router has added the new route to its multicast routing table, it then accepts multicast packets for these pairs on any interface in the virtual router and forwards them accordingly. Multicast routes established before you issue this command are not affected.

Specifying Unicast Routes for RPF

You can specify that IS-IS, OSPF, or RIP routes be available for RPF. Routes available for RPF appear in the multicast view of the routing table.

```
host1(config)#router ospf 1
host1(config-router)#ip route-type multicast
```

Defining Permanent IP Multicast Forwarding Entries

An mroute is a multicast traffic flow (a (Source, Group) entry used for forwarding multicast traffic). By default, forwarding mroutes (with a valid RPF incoming interface) are timed out if data for them is not received for 210 seconds. However, you can specify an mroute as permanent by using the **ip multicast-routing permanent-mroute** command.

```
host1(config)#ip multicast-routing permanent-mroute routes1
```

You can use the **no** version to prevent any new mroutes from becoming permanent. To remove existing permanent mroutes, use the **clear ip mroute** command.

Defining a Multicast Bandwidth Map

Multicast interface-level admission control, port-level admission control, and QoS adjustment all use a single multicast bandwidth map. The multicast bandwidth map is a route map that uses the **set admission-bandwidth**, **set qos-bandwidth**, **set admission-bandwidth adaptive**, or **set qos-bandwidth adaptive** commands. The

adaptive commands configure an autosense mechanism for measuring the multicast bandwidth.

NOTE: Even though you can include any of the preceding commands several times in a route map entry, only the last **admission-bandwidth** command or **qos-bandwidth** command in the bandwidth map is used. In other words, if you included the **set qos-bandwidth** command first and then the **set qos-bandwidth adaptive** command, the bandwidth map uses the **set qos-bandwidth adaptive** command.

Interface-level and port-level admission control is performed when an OIF on the interface or port is added to the mroute for a given (S,G) multicast data stream and the multicast bandwidth map contains a **set admission-bandwidth** or **set admission-bandwidth adaptive** action for that (S,G).

QoS adjustment is performed on the joining interface when an OIF is added to the mroute for a given (S,G) data stream and the multicast bandwidth map contains a **set qos-bandwidth** or **set qos-bandwidth adaptive** action for that (S,G).

You can prioritize the traffic by configuring a priority value for the <S, G> data stream on a physical port by issuing the **set priority** command. Dynamic multicast admission control enables only prioritized groups to join the interface after the configured priority limit is reached on the physical port. The system records the priority when a new <S, G> entry is created.

NOTE: You can create a single route map with the **set admission-bandwidth** command, the **set qos-bandwidth** command, or both. However, creating an entry with only one of these **set** commands enables only that specific function for the matched address (that is, only multicast traffic admission control or only QoS adjustment). The same is true for the **adaptive** commands.

Autosense Mechanism Overview

Video bandwidth is typically considered to be a constant rate—2 Mbps for standard definition television (SDTV) and 10 Mbps for high definition television (HDTV). However, in reality, and depending on achievable video compression, the bit rate can vary. For example, HDTV streams (using MPEG4 or WM9 encoding) can vary between 6 Mbps (for low-action programs) to 10 Mbps (for a fast-paced, high-action programs). The autosense mechanism causes the bandwidth value, used for admission control and QoS adjustment, to be the actual measured rate of the stream. Using this feature to measure the actual bandwidth avoids the need to configure arbitrary bandwidth limits and enables a channel to be reassigned to a different (S, G) without requiring a bandwidth map to be changed.

Related Documentation

- Adaptive Mode Mechanism Overview on page 12
- Defining a Multicast Bandwidth Map on page 8

Example: Configuring a Multicast Bandwidth Map

The following example creates a multicast bandwidth map for both multicast traffic admission control and QoS adjustment:

- Requirements on page 10
- Overview on page 10
- Configuring an IPv4 Multicast Bandwidth Map on page 10

Requirements

This example uses the following hardware and software components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Before you begin configuring multicast on IPv4 interfaces, you must:

- Configure IPv4 interfaces. For more information about configuring IPv4 interfaces, see *Configuring IPv4* in *JunosE IP, IPv6, and IGP Configuration Guide*.

Overview

The multicast bandwidth map is a route map that uses the **set admission-bandwidth**, **set qos-bandwidth**, **set admission-bandwidth adaptive**, or **set qos-bandwidth adaptive** commands. Multicast interface-level admission control, port-level admission control, and QoS adjustment all use a single multicast bandwidth map.

The **adaptive** commands configure an auto-sense mechanism for measuring the multicast bandwidth.

Configuring an IPv4 Multicast Bandwidth Map

Configuring a Route Map

Step-by-Step Procedure

Define a route map using the **set admission-bandwidth** and **set qos-bandwidth** commands.

NOTE: In this example, you can replace the **set admission-bandwidth** command and **set qos-bandwidth** command with their **adaptive** command counterparts.

1. Define a route map.

   ```
   [edit]
   host1(config)#route-map mcast-bandwidths permit 10
   ```

2. Match the route map to an access list.

   [edit]
   host1(config-route-map)#**match ip address sdtv**

3. Configure multicast bandwidths.

   [edit multicast bandwidths for admission control and QoS adjustment]
   host1(config-route-map)#**set admission-bandwidth 2000000**
   host1(config-route-map)#**set qos-bandwidth 2000000**

4. Configure the route map.

   [edit]
   host1(config-route-map)#**route-map mcast-bandwidths permit 20**

5. Match the route map to an access list.

   [edit]
   host1(config-route-map)#**match ip address hdtv**

6. Configure multicast bandwidths.

   [edit multicast bandwidths for admission control and QoS adjustment]
   host1(config-route-map)#**set admission-bandwidth 10000000**
   host1(config-route-map)#**set qos-bandwidth 10000000**

Configuring an Access List

Step-by-Step Procedure

Define the access list for use by the **match ip address** command to match (S,G) and (*,G) entries.

NOTE: You can also define a prefix-list or a prefix-tree for use by the **match ip address** command to match (S,G) and (*,G) entries.

1. Configure access lists.

   [edit access lists]
   host1(config)#**access-list sdtv permit ip host 31::1 ff3e::0/112**
   host1(config)#**access-list hdtv permit ip host 32::1 ff3e::0/112**
   host1(config)#**access-list hdtv permit ip host 32::2 ff3e::0/112**

Related Documentation

- Defining a Multicast Bandwidth Map on page 8
- access-list
- match ip address
- set admission-bandwidth
- set qos-bandwidth

Adaptive Mode Mechanism Overview

You configure the auto-sense mechanism in the multicast bandwidth using the **set admission-bandwidth adaptive** command, **set qos-bandwidth adaptive** command, or both. For example:

```
host1(config)#route-map mcast-bandwidths permit 10
host1(config-route-map)#match ip address sdtv
host1(config-route-map)#set admission-bandwidth adaptive
host1(config-route-map)#set qos-bandwidth adaptive
host1(config-route-map)#end
```

In this example, any stream with an (S,G) that matches the sdtv access list performs adaptive bandwidth detection for admission control and QoS adjustment.

A rate measurement mechanism runs on the ingress line card that polls the forwarding controller (FC) to obtain statistics for each mroute. This mechanism then reports the rate measurement to the SRP to update the bandwidth map. By computing the average bandwidth over a relatively short sampling period (T1; 5 seconds), the measurement approximates the peak bandwidth of the multicast stream.

As an example, assume that a new mroute (S1, G1) is added to the interface controller (IC) at time t0.

Figure 1: Example of Adaptive IPv4 Multicast Bandwidth Detection

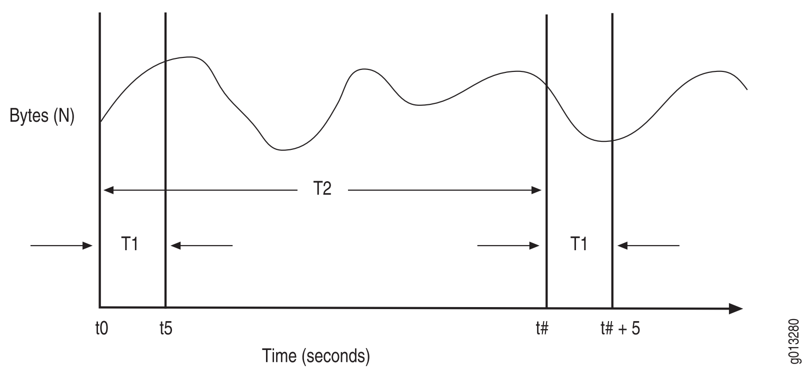

To calculate the measured bandwidth of a stream, the router uses the following equation:

$R = (N_{t+5} - N_t) / 5$

Where

R = Calculated bandwidth of the stream during each sampling interval

N_t = Bytes measured at the start of each sampling period (t seconds)

N_{t+5} = Bytes measured at the end of each sampling period (t+5 seconds)

NOTE: When the mroute is first installed in the FC (at t = 0), R_0 is undetermined. For multicast admission control no joins are admitted until the first bandwidth measurement is computed (that is, for admission control, R0 is considered to be infinite). Similarly, no QoS adjustment occurs until the first bandwidth measurement is computed (that is, for QoS adjustment, R0 is considered to be zero [0]).

Using the previous graph as a reference, the first bandwidth rate ($R1_0$) and at time t_5 (N_5) and the bytes received values are subtracted and divided by the sampling period T_1 to yield the average rate. This process is repeated every sampling interval, T_2, to yield rates R1, R2, R3, and so on.

The first two sampling interval calculations are as follows:

$R_1 = (N_5 - N_0)/5$

$R_2 = (N_{\#+5} - N_{\#})/5$

The router maintains a history of bandwidth measurements (H) for each mroute, up to a maximum of M measurements. The actual rate, R, reported to the SRP is the maximum rate measured in those H samples.

To minimize the IC to SRP traffic generated by the rate measurements, the IC reports a bandwidth change only when a newly computed rate (R#) differs from the current rate by a specified threshold. When R_s is computed at time t = 5 seconds, R is set to R_1. A rate update occurs whenever a newly calculated rate (R) differs from R_1 by at least a threshold value (specified as a percentage, P) of the measured peak bandwidth. This calculation is as follows:

R = Rt, if and only if the absolute value of (R - Rt) > P * R.

Table 4 on page 13 lists values assigned to variables associated with this algorithm.

Table 4: Adaptive Mode Algorithm Values

Variable	Value	Units	Description
T1	5	Seconds	Sampling period; the time in which a sample is taken
T2	0	Seconds	Sampling interval; zero (0) seconds indicates continuous sampling
H	12	Samples	Number of history samples over which to compute measurement
M	12	Samples	Maximum number of samples maintained in history
P	1	Percent	Threshold value; percent difference by which a newly calculated rate must differ from the measured peak bandwidth before a rate update occurs

Related Documentation

- Autosense Mechanism Overview on page 9
- Defining a Multicast Bandwidth Map on page 8
- match ip address
- route-map
- set admission-bandwidth
- set qos-bandwidth

Multicast QoS Adjustment for IPv4

When the router uses multicast OIF mapping, any multicast streams that a subscriber receives bypass any configured QoS treatment for that subscriber interface. The Multicast QoS adjust feature provides a way in which the router can account for this multicast traffic.

The following topics provide two possible configuration cases for using multicast QoS adjustment.

Multicast OIF Mapping Case

Multicast OIF mapping enables the router to decrease the inefficiencies associated with replicating streams of multicast traffic. Using OIF maps, IGMP joins that the router receives on a subscriber interface can be mapped to a special interface for forwarding. This special interface can be on a different physical port or line module from that of the join interface.

Using this mapping function, the router can send a single copy of each multicast stream over the special interface and the access nodes are configured to perform any final replication to the subscribers and merge unicast and multicast data flows onto the subscriber interfaces as necessary. See Figure 2 on page 15.

Figure 2: Multicast OIF Mapping

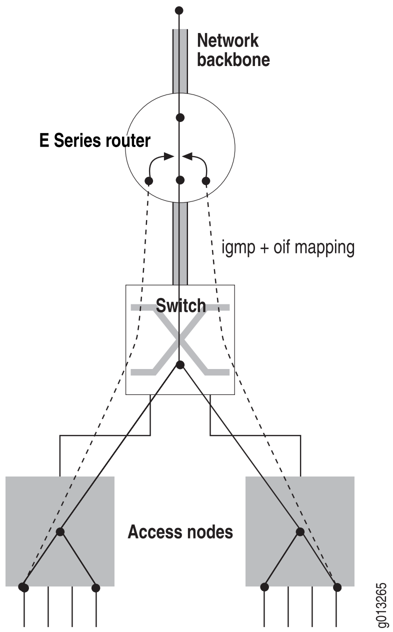

One disadvantage to using multicast OIF mapping is that the multicast traffic bypasses any QoS treatment that is applied to subscriber interfaces. Configuring QoS adjustment resolves this problem. With QoS adjustment configured, when a subscriber requests to receive a multicast stream (or, more appropriately, when an OIF is added to the mroute), the router reduces the unicast QoS bandwidth applied to the subscriber interface (that is, the join interface) by the amount of bandwidth for that multicast stream.

Related Documentation

- Configuring Group Outgoing Interface Mapping on page 62
- Multicast Traffic Receipt Without Forwarding on page 15
- Activating Multicast QoS Adjustment Functions on page 16

Multicast Traffic Receipt Without Forwarding

In this case, the router is not given the responsibility of forwarding multicast streams. Instead, the service provider arranges for the router to receive the multicast streams so the router can detect the flow and perform QoS adjustment. An OIF map is installed that maps the traffic streams to a loopback interface configured for IGMP version passive. This means that when the traffic is received, a null mroute is installed (that is, an mroute with an empty OIF list) and the router applies the QoS adjustment to the join interface. See Figure 3 on page 16.

NOTE: Ensure that PIM-SM (or any other upstream multicast protocol) is informed of the group (or source-group) interest.

Figure 3: Multicast Traffic Receipt Without Forwarding

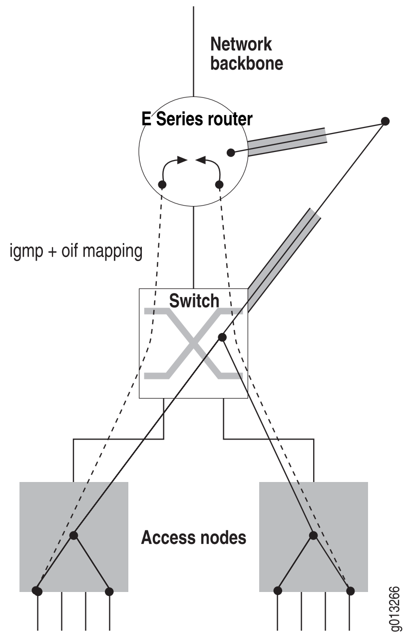

Related Documentation

- Multicast OIF Mapping Case on page 14
- Activating Multicast QoS Adjustment Functions on page 16

Activating Multicast QoS Adjustment Functions

The **ip multicast-routing bandwidth-map** command activates the specified bandwidth map. By activating the bandwidth map, this command also activates the multicast QoS adjustment function contained in the bandwidth map.

CAUTION: To activate multicast QoS adjustment, you must first create a bandwidth map. See "Defining a Multicast Bandwidth Map" on page 8 for details.

To enable the QoS adjust function on the router with the configured route map:

- Issue the **ip multicast-routing bandwidth-map** command in the Global Configuration mode:

  host1(config)#**ip multicast-routing bandwidth-map mcast-bandwidths**

You can use the **no** version to disable the multicast QoS adjustment function on the router.

Related Documentation

- Multicast OIF Mapping Case on page 14
- Multicast Traffic Receipt Without Forwarding on page 15
- Monitoring the IP Multicast Status on a Virtual Router on page 52

- ip multicast-routing bandwidth-map

Related Documentation

- Configuring Group Outgoing Interface Mapping on page 62
- IP Multicast Bandwidth Adjustment for QoS Overview
- Parameter Definition Attributes for QoS Administrators Overview
- ip multicast-routing bandwidth-map

Hardware Multicast Packet Replication Overview

You can configure IPv4 multicast to replicate packets to optimized hardware on a logical port instead of using the forwarding controller (FC) on the router.

The bandwidth between the line module and the I/O module or IOA on the E Series router is limited. A high-density Ethernet module provides eight physical ports that can consume the bandwidth between the line module and the I/O module or IOA before providing enough traffic to support egress line rate for all of these ports.

Figure 4 on page 19 displays how multicast traffic is typically replicated on the line module. Each of these replicated packets is transmitted from the line module to the I/O module or IOA.

Figure 4: Packet Flow Without Hardware Multicast Packet Replication

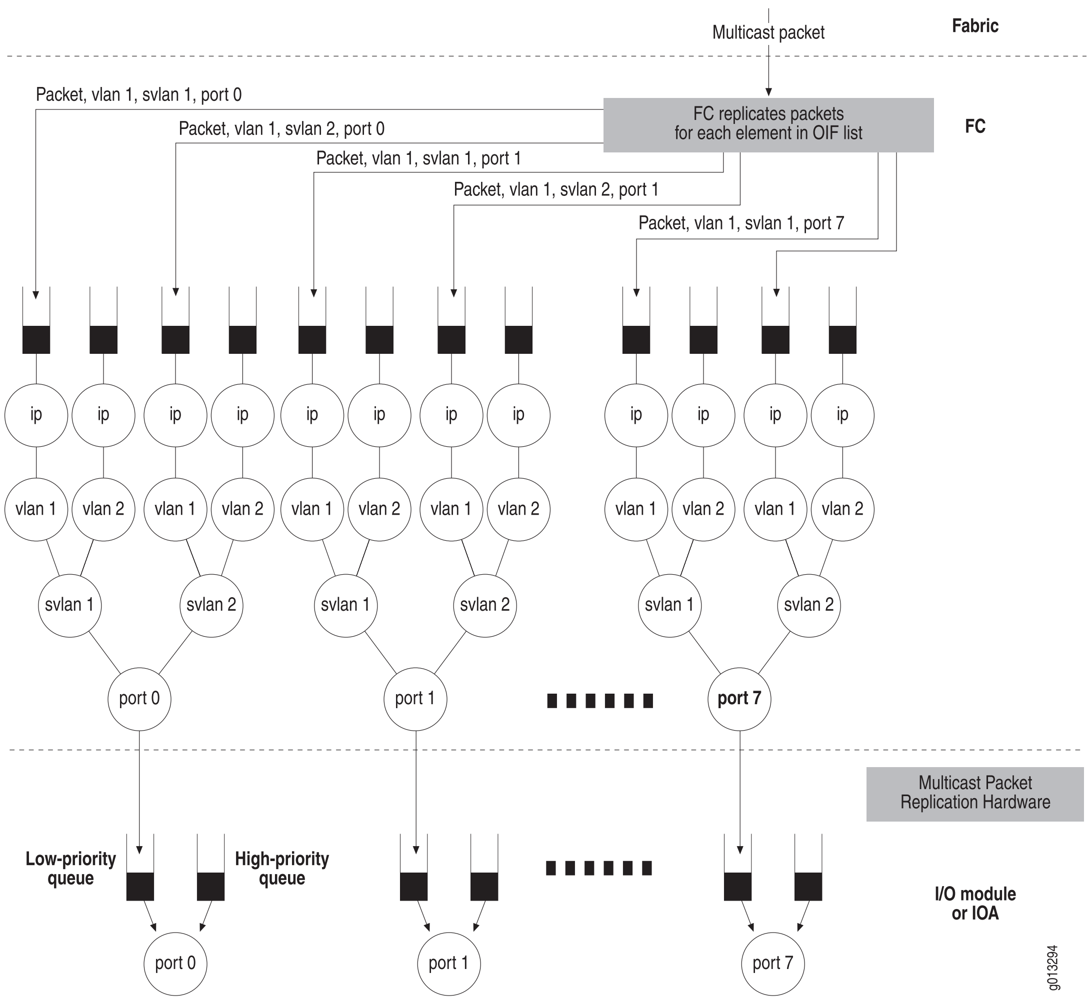

The hardware multicast packet replication feature enables you to configure multicast traffic for a VLAN or S-VLAN to be replicated on the I/O module or IOA so that only one copy of the packet is transmitted from the line module to the I/O module or IOA. Replication for each of the ports is performed on the I/O module or IOA.

Configuring hardware multicast packet replication for high-density Ethernet is useful when you want to provide the same multicast stream out of some or all of the ports, such as for IP television (IPTV). Configuring hardware multicast packet replication enables you to:

- Reduce the number of packets sent from the FC to the module.
- Reduce the CPU consumed by the FC processing each elaboration of the packet.

You can use the additional bandwidth to increase the bandwidth of multicast traffic out of each of the Gigabit Ethernet ports.

Figure 5 on page 20 displays the flow of a multicast packet using the hardware multicast packet feature.

Figure 5: Packet Flow with Hardware Multicast Packet Replication

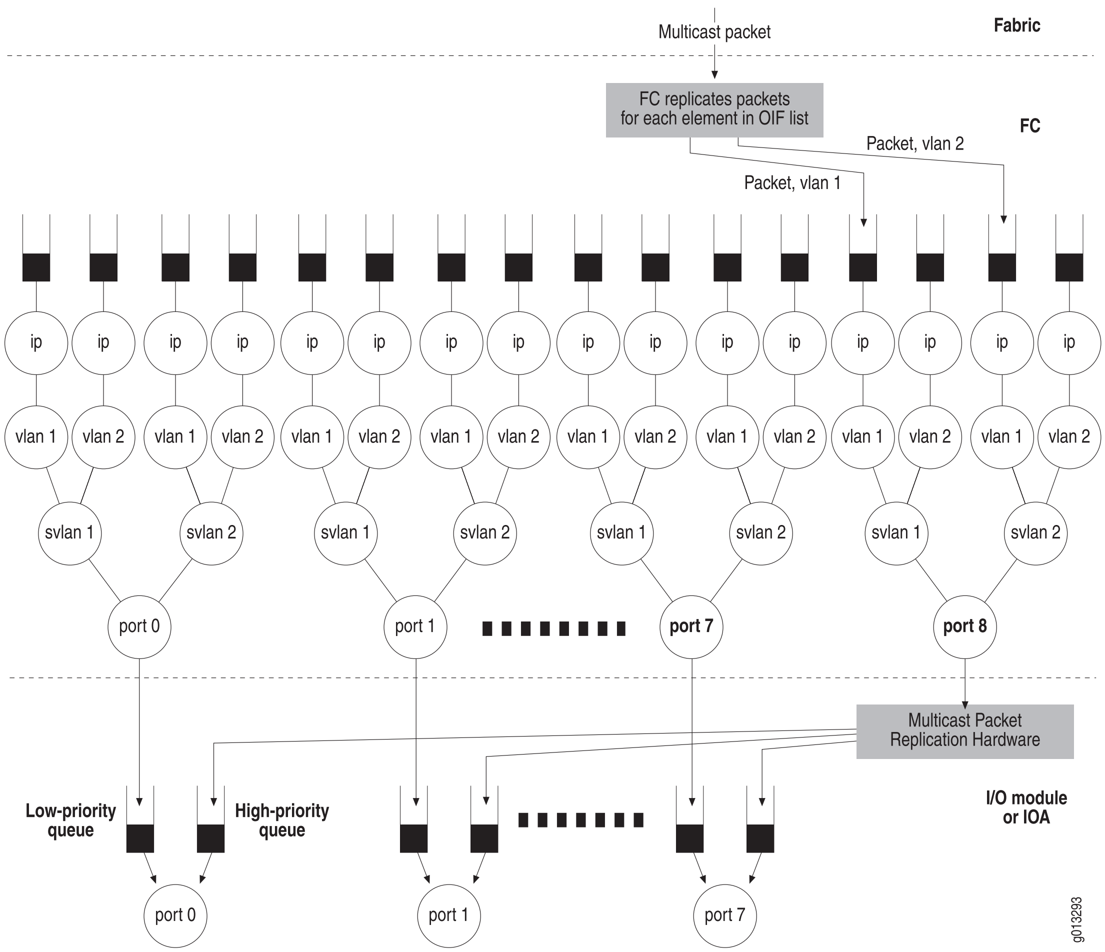

Each high-density Ethernet module has eight physical ports, numbered 0–7. A logical port is available for the hardware multicast packet replication feature, numbered port 8.

JunosE tracks the OIFs in an mroute that have been redirected to use the hardware multicast packet replication hardware. The system accepts only egress multicast traffic to traverse the interface stack on the enabled port. The system drops unicast traffic that is routed to this port.

Each port on the I/O module or IOA displayed in Figure 5 on page 20 has two queues. These queues are further down the egress path than the queues found on the line module and populated by the FC.

The low-priority queue is dedicated to packets that are received from the line module queues that are dedicated to the physical ports. This queue blocks when full and provides backpressure to the line module. This queue services unicast and multicast traffic that is not using the hardware multicast packet replication feature.

The high-priority queue is dedicated to packets that are received from the line module queue for port 8. This queue is serviced at a higher priority than the first queue, and drops packets when full.

For more information about high-density Ethernet, see *Configuring Ethernet Interfaces* in the *JunosE Physical Layer Configuration Guide*.

This topic discusses the following:

Supported Modules and Encapsulations

You can enable hardware multicast packet replication on port 8 of the following high-density Ethernet modules:

- GE-8 I/O module (pairs with the GE-HDE line module)
- ES2-S1 GE-8 IOA (pairs with the ES2 4G LM and the ES2 10G LM)

When enabled, the hardware multicast packet replication feature defines the encapsulation of the egress multicast packet. The following encapsulations are supported:

- IPv4 over Gigabit Ethernet
- IPv4 over VLAN
- IPv4 over S-VLAN

NOTE: **802.3ad link aggregation group (LAG) bundles do not support hardware multicast packet replication.**

The hardware multicast packet replication feature also provides an interface over which you can configure the following:

- IP MTU
- Ethernet MTU

- Egress IP policy
- Egress VLAN policy
- QoS

Relationship with OIF Mapping

Multicast OIF mapping enables the router to decrease the inefficiencies associated with replicating streams of multicast traffic. Using OIF maps, IGMP joins that the router receives on a subscriber interface can be mapped to a dedicated multicast VLAN.

The hardware multicast packet replication feature enables you to redirect each of the IP interfaces on a line module over a dedicated multicast VLAN to a single IP interface over port 8. The FC is only required to send a single packet per dedicated multicast VLAN to the I/O module or IOA. The module then replicates this packet to the appropriate ports.

Related Documentation

- Configuring Group Outgoing Interface Mapping on page 62
- Configuring Hardware Multicast Packet Replication on page 24
- Hardware Multicast Packet Replication Considerations on page 22

Hardware Multicast Packet Replication Considerations

When configuring hardware multicast packet replication, the following considerations apply.

- Do not configure or transmit routing protocols over port 8. The FC drops traffic routed to an IP interface stacked over port 8.
- We recommend that you configure the IP address of the IP interface over port 8 to be unnumbered.
- You must configure the same VLAN settings over the port (logical port 8, in this case) on which you configure the multicast replication feature for I/O modules or IOA and the physical ports.

  You cannot create the following configurations:

  - When two IP interfaces configured over a port reference the same IP interface over port 8. The system does not accept this configuration attempt because you typically configure the hardware multicast packet replication feature to redirect multicast traffic over one VLAN, then redirect it to the same VLAN on port 8.
  - When the IP interface configured with the hardware multicast packet replication attribute is not installed on a line module that supports hardware multicast packet replication.

  - When the IP interface designated by the hardware multicast packet replication attribute is not installed on a line module that supports hardware multicast packet replication.
  - When the IP interface designated by the hardware multicast packet replication attribute is not on the same line module as the IP interface configured with this attribute.

- When you configure a unique source MAC address for VLANs on port 8, the hardware multicast packet replication hardware stamps the source MAC address on the VLAN, overwriting any MAC address that you configured. For more information, see *Configuring Ethernet Interfaces* in the *JunosE Physical Layer Configuration Guide*.
- The regular multicast implementation utilizes interface stacking that provides a unique IP attachment point for each elaboration of the egress multicast packet.

  For the hardware multicast packet replication feature, you must attach policies to an interface stack over port 8 that defines the encapsulation of the egress multicast traffic. The system supports policies over port 8 just as it is above any of the other ports on this line module.

  Policies applied to the interface stack over port 8 affect the packets traversing this stack whether or not the packet is destined for one port or all of the physical ports. Therefore, you cannot apply different egress policies to multicast traffic for the interfaces stacked above different ports, or rate limit on an individual interface over a port. You also cannot monitor policy statistics on individual interfaces over a port.

  Instead, you can apply egress policy to an interface stacked over port 8. The system applies the policy before the packet has been elaborated for each of the ports.
- The JunosE QoS component provides hierarchical egress scheduling and shaping on Gigabit Ethernet ports 0–7. The regular multicast implementation replicates packets on the FC, with each replicated packet placed on a line module queue destined for a single physical port. The line module queue can also receive QoS behavior specific to that queue.

  For the hardware multicast packet replication feature, the FC does not replicate the packet for each of the individual ports. Instead, it places the packet on a special queue destined for port 8.

  You can configure QoS on the packets flowing through port 8, but this has limited value because each packet passed through this port can be transmitted through one of more of the physical ports. Therefore, the packets placed on this special queue might not receive the same QoS behavior as ports 0–7.

  We recommend that you configure the network so the I/O or IOA queues are not oversubscribed. The traffic transmitted by the physical port is a combination of packets from the two I/O or IOA queues. When the sum of the packets in these queues is greater than line rate, the system can drop traffic that is not using hardware multicast packet replication.

  When you configure a traffic shaper on a physical port and configure hardware multicast packet replication, the packets created using the feature avoid the traffic shaper for that port. To control this, you can use traffic shaper on the physical port and port 8. The sum of the traffic shapers must be less than or equal to the line rate of the port.

A traffic shaper on port 8 can result in the overall utilization of egress bandwidth for any one port being less the line rate because the packets being replicated might not be transmitted to every port. Packets destined to some of the ports contribute to the traffic shaping for all of the ports on the I/O module or IOA.

Related Documentation

- Hardware Multicast Packet Replication Overview on page 18
- Configuring Hardware Multicast Packet Replication on page 24

Configuring Hardware Multicast Packet Replication

The following topics explain how to configure hardware multicast packet replication without and with OIF-mapping:

- Configuring Hardware Multicast Packet Replication Without OIF-Mapping on page 24
- Configuring Hardware Multicast Packet Replication With OIF-Mapping on page 25

Configuring Hardware Multicast Packet Replication Without OIF-Mapping

To configure hardware multicast packet replication without OIF-mapping:

1. Configure port 8 on a high-density Ethernet module to accept redirected egress multicast traffic.
   a. Specify the Gigabit Ethernet interface on port 8.
   b. Create a VLAN major interface.
   c. Create a VLAN subinterface.
   d. Assign a VLAN ID.
   e. Configure an unnumbered IP interface.
   f. Enable IGMP on the interface with only multicast-data-forwarding capability.

   ```
   host1(config)#interface gigabitEthernet 2/8
   host1(config-if)#encapsulation vlan
   host1(config-if)#interface gigabitEthernet 2/8.1
   host1(config-if)#vlan id 1
   host1(config-if)#ip unnumbered loopback 0
   host1(config-if)#ip igmp version passive
   ```

2. Configure an IP interface to redirect egress multicast traffic to port 8.
   a. Create a VLAN subinterface.
   b. Assign a VLAN ID.
   c. Assign an IP address.
   d. Configure the interface to redirect egress multicast traffic to port 8.

```
host1(config)#interface gigabitEthernet 2/0.101
host1(config-if)#vlan id 1
host1(config-if)#ip address 10.1.1.1 255.255.255.0
host1(config-if)#ip multicast ioa-packet-replication gigabitEthernet 2/8.1
```

Configuring Hardware Multicast Packet Replication With OIF-Mapping

This section describes how to configure hardware multicast packet replication with OIF-mapping.

1. Configure port 8 on a supported high-density Ethernet module to accept redirected egress multicast traffic. For information about supported high-density Ethernet modules see "Hardware Multicast Packet Replication Overview" on page 18.
2. Use OIF maps to map the subscriber IGMP interfaces (C-VLANs) to the dedicated multicast VLAN (M-VLAN). The dedicated M-VLAN should be located on the line module containing the IOA replication interface. The C-VLAN and M-VLAN can either be on the same or different line modules.
3. Configure the dedicated M-VLAN to redirect egress multicast traffic to port 8.

Related Documentation

- Hardware Multicast Packet Replication Overview on page 18
- Monitoring Multicast Routes When OIF Mapping Is Configured on page 43

- encapsulation vlan
- ip igmp version
- ip multicast ioa-packet-replication
- ip unnumbered

Controlling Multicast Traffic on the Interface

You can control multicast traffic on the interface with the following topics:

- Blocking Mroutes on page 25
- Enabling Interface Admission Bandwidth Limitation on page 26
- OIF Interface Reevaluation on page 27

Blocking Mroutes

By default, when an interface that is configured with one or more multicast protocols (for example, PIM or IGMP) receives multicast traffic, even when the scope of that traffic exceeds link-local, the virtual router creates an mroute. You can use the **ip block-multicast-sources** command to block all multicast traffic with a scope larger than link-local (for example, global) and prevent mroute creation under these conditions.

NOTE: **Issuing this command does not affect reception of link-local multicast packets.**

To prevent mroute creation by blocking multicast traffic that has a scope larger than link-local (for example, global):

- Issue the **ip block-multicast-sources** command in Interface Configuration mode:

  host1(config-if)#**ip block-multicast-sources**

You can use the **no** version to restore the default behavior of creating mroutes on received multicast packets.

Related Documentation

- Monitoring Multicast Forwarding Entries on page 37
- Monitoring Active Multicast Routes on page 40
- Monitoring Multicast Entries in a Source or Group on page 43
- Monitoring Multicast Routes When OIF Mapping Is Configured on page 43
- Monitoring Multicast Statistics on page 46
- Monitoring Summary Information of Multicast Routes on page 49
- Monitoring Multicast Protocols Enabled on the Router on page 50

- ip block-multicast-sources

Enabling Interface Admission Bandwidth Limitation

Interface-level multicast admission control is performed when an OIF on the interface is added to the mroute for a given (S,G) multicast data stream and the multicast bandwidth map contains a **set admission-bandwidth** action for that (S,G). When enabled, the admission-bandwidth for a particular (S,G) is read from the multicast bandwidth map and recorded in the mroute when the (S,G) mroute is created.

CAUTION: **Before you can limit interface-level admission bandwidth, you must first create a bandwidth map. See** "Defining a Multicast Bandwidth Map" on page 8 **for details.**

When an OIF is subsequently added to the mroute, the OIF is blocked from forwarding data if the additional bandwidth contributed by the (S,G) exceeds the admission-bandwidth limit for the interface. In JunosE releases earlier than Release 12.0.x, in an OIF mapping scenario where the DSLAM does not perform per-subscriber multicast admission control, the router disregards the multicast admission bandwidth limit configured on the join interface. If the limit configured on the mapped interface exceeds the admission-bandwidth limit for the interface, the router blocks the mapped interface from forwarding data.

Now, in an OIF mapping scenario where the DSLAM performs per-subscriber multicast admission control, the router checks the bandwidth limit configured on the join interface. If the multicast stream is forwarded over the mapped interface, the router admits the multicast stream and forwards the stream to the join interfaces whose bandwidth does not exceed the configured bandwidth limit. The router also performs QoS Adjust for the multicast stream on the unblocked (forwarding) subscriber interfaces. The router does not replicate the stream to the subscriber interfaces whose bandwidth exceeds the configured bandwidth limit and it does not perform QoS Adjust for the multicast stream on the blocked subscriber interfaces.

If the multicast stream is not forwarded over the mapped interface, the router blocks the multicast stream and does not forward the stream. The router also does not perform QoS Adjust for the multicast stream on the blocked subscriber interface.

You can use the **ip multicast admission-bandwidth-limit** command to enable multicast admission control on interfaces (including dynamic IP interfaces) that are configured to run IGMP. You can also use this command on a PIM (sparse-mode, dense-mode, or sparse-dense-mode) interface if IGMP is configured on the interface (including the **ip igmp version passive** command).

To limit the bandwidth for an interface that accepts IGMP groups:

- Issue the **ip multicast admission-bandwidth-limit** command in Interface Configuration mode.

  host1(config-if)#**ip multicast admission-bandwidth-limit 2000000**

You can use the **no** version to remove the bandwidth limitation for the interface.

Related Documentation

- OIF Interface Reevaluation on page 27
- Monitoring Multicast Forwarding Entries on page 37
- Monitoring Active Multicast Routes on page 40
- Monitoring Multicast Routes When OIF Mapping Is Configured on page 43
- Monitoring Multicast Statistics on page 46
- Monitoring Multicast Protocols Enabled on the Router on page 50

- ip igmp version
- ip multicast admission-bandwidth-limit
- set admission-bandwidth

OIF Interface Reevaluation

If you change the admission bandwidth for an interface, all mroutes with that interface as an OIF are reevaluated as follows:

- If the bandwidth limit is increased, blocked OIFs may become unblocked. If the interface is a blocked OIF on multiple mroutes, the order in which the mroutes are visited, and which (S,G) streams become unblocked, is not specified.
- If the bandwidth limit is decreased, no currently admitted OIFs are blocked. However, no new OIFs are admitted until the total admitted bandwidth for the interface drops below the new limit.
- If the bandwidth is increased to the point that the bandwidth limit for an interface is now exceeded, no currently admitted OIFs for the affected mroutes are blocked.

However, no new OIFs are admitted until the total admitted bandwidth drops below the configured limit.

NOTE: **If the multicast bandwidth map that includes the set admission-bandwidth command is changed, all affected mroutes are reevaluated in the same manner described previously.**

As an example of this function, if the interface has accepted a total bandwidth of 2000000 bps, and you set a limit of 1000000 bps on the interface, the router does not disconnect any already connected OIFs but prevents the interfaces from accepting any more groups. Over time, some groups leave the interfaces and, eventually, the interface limit of 1000000 bps is reached and maintained by the router.

If you set limits for both a port and interfaces on that port, the router uses the lower of the two limits when determining whether or not an interface can accept any new IGMP groups. For example, if you specify an admission bandwidth limit of 2000000 bps for the port and 3000000 bps groups for each interface, additional groups can only be accepted until the port limit of 2000000 bps is reached.

Related Documentation

- Defining a Multicast Bandwidth Map on page 8
- Enabling Interface Admission Bandwidth Limitation on page 26
- OIF Port Reevaluation on page 31
- set admission-bandwidth

Related Documentation

- Creating Mroute Port Limits on page 29
- Enabling Port-Level Admission Bandwidth Control on page 29
- Dynamic Port Admission Bandwidth Control on page 30
- OIF Port Reevaluation on page 31
- ip block-multicast-sources
- ip igmp version
- ip multicast admission-bandwidth-limit
- set admission-bandwidth

Controlling Mutlicast Traffic on the Port

You can control multicast traffic on the port with the following topics:

Creating Mroute Port Limits

When a multicast forwarding entry (that is, an mroute) is added with an outgoing interface (OIF) on a port, the OIF count for that port is incremented. If you configure a port limit, and the OIF count on the port exceeds that limit, no OIFs on that port are added to mroutes (that is, OIFs are blocked).

To configure a limit on the number of mroute OIFs that can be added across different virtual routers on a port:

- Issue the **mroute port limit** command in Global Configuration mode.

```
host1(config)#mroute port 3/0 limit 10
```

You can use the **no mroute port limit** command to remove any OIF port limits.

Related Documentation

- Monitoring Multicast Forwarding Entries on page 37
- Monitoring Active Multicast Routes on page 40
- Monitoring Multicast Routes When OIF Mapping Is Configured on page 43
- Monitoring Multicast Statistics on page 46
- Monitoring Summary Information of Multicast Routes on page 49
- Monitoring Multicast Routes on Virtual Router Ports on page 53

- mroute port limit

Enabling Port-Level Admission Bandwidth Control

Port-level multicast admission control is performed when an OIF on that port is added to the mroute for a given (S,G) multicast data stream and the multicast bandwidth map contains a **set admission-bandwidth** action for that (S,G).

When enabled, the admission-bandwidth for a particular (S,G) is read from the multicast bandwidth map and recorded in the mroute when the (S,G) mroute is created. If you configure a port limit and the OIF count on the port exceeds that limit, no OIFs on that port are added to mroutes (that is, OIFs are blocked).

When a multicast forwarding entry (an mroute) is added with an outgoing interface, OIF is blocked from forwarding data if the additional bandwidth contributed by the (S,G) would exceed the admission-bandwidth limit for the port on which the interface resides.

CAUTION: Before you can limit port-level admission bandwidth, you must first create a bandwidth map. See "Defining a Multicast Bandwidth Map" on page 8 for details.

You can use the **mroute port admission-bandwidth-limit** command to limit the total multicast bandwidth that can be admitted on a port. The admitted bandwidth is summed across all virtual routers with IPv4 and IPv6 mroutes that have OIFs on the port.

NOTE: Admission bandwidth values for a given (S,G) mroute are determined from the bandwidth map. See "Defining a Multicast Bandwidth Map" on page 8 for details.

Related Documentation

- Dynamic Port Admission Bandwidth Control on page 30
- OIF Port Reevaluation on page 31
- Monitoring Multicast Forwarding Entries on page 37
- Monitoring Active Multicast Routes on page 40
- Monitoring Multicast Routes When OIF Mapping Is Configured on page 43
- Monitoring Multicast Statistics on page 46
- Monitoring Multicast Protocols Enabled on the Router on page 50
- Monitoring Multicast Routes on Virtual Router Ports on page 53

- mroute port admission-bandwidth-limit
- set admission-bandwidth

Dynamic Port Admission Bandwidth Control

You can configure the system to dynamically limit the total multicast bandwidth that can be admitted on a port. The system performs dynamic port-level admission control when an OIF on that port is added to the mroute for a given <S, G> multicast stream.

After the priority bandwidth limit on the port is reached, OIFs on the prioritized <S, G> are only allowed to forward the traffic and unprioritized <S, G> streams are blocked from forwarding data on the OIF.

To enable a priority value for the <S, G> multicast stream, issue the **set priority** command in the multicast bandwidth map. A priority value of 0 indicates an unprioritized stream and any value other than 0 indicates a prioritized stream. Currently there is no support for classification of prioritized streams.

```
host1(config)#mroute port admission-bandwidth-limit 3000000
```

You can configure limits for the bandwidth that is dynamically admitted on the port. The priority bandwidth limit controls the priority bandwidth admitted on a port. The hysteresis limit sets the minimum priority bandwidth limit before the system evaluates mroutes and admits any blocked OIFs.

Related Documentation

- Defining a Multicast Bandwidth Map on page 8
- Enabling Port-Level Admission Bandwidth Control on page 29

- OIF Port Reevaluation on page 31
- mroute port admission-bandwidth-limit
- set priority

OIF Port Reevaluation

If you change the admission bandwidth for a port, all mroutes with an OIF on that port are reevaluated as follows:

- If the bandwidth limit is increased, blocked OIFs can become unblocked. However, the order in which the mroutes are visited, and which (S,G) streams become unblocked, is not specified.
- If the bandwidth limit of a port is decreased, no currently admitted OIFs are blocked. However, no new OIFs are admitted until the total admitted bandwidth for the port drops below the new limit.
- If the bandwidth is increased to the point that the bandwidth limit for an interface is now exceeded, no currently admitted OIFs for the affected mroutes are blocked. However, no new OIFs are admitted until the total admitted bandwidth drops below the configured limit.

NOTE: **If the multicast bandwidth map that includes the set admission-bandwidth command is changed, all affected mroutes are reevaluated in the same manner described previously.**

As an example of this function, if the port has accepted a total bandwidth of 3000000 bps, and you set a limit of 2000000 bps on the port, the router does not disconnect any already connected OIFs but prevents the interfaces from accepting any more groups. Over time, some groups leave the interfaces and, eventually, the port limit of 2000000 bps is reached and maintained by the router.

If you set limits for both a port and interfaces on that port, the router uses the lower of the two limits when determining whether or not an interface can accept any new IGMP groups. For example, if you specify an admission bandwidth limit of 2000000 bps for the port and 3000000 bps groups for each interface, additional groups can only be accepted until the port limit of 2000000 bps is reached.

Related Documentation

- OIF Interface Reevaluation on page 27
- Defining a Multicast Bandwidth Map on page 8
- Enabling Port-Level Admission Bandwidth Control on page 29
- Dynamic Port Admission Bandwidth Control on page 30
- set admission-bandwidth

Related Documentation

- Blocking Mroutes on page 25
- Enabling Interface Admission Bandwidth Limitation on page 26
- OIF Interface Reevaluation on page 27
- mroute port admission-bandwidth-limit
- mroute port limit
- set admission-bandwidth

Deleting Multicast Forwarding Entries

You can clear one or more forwarding entries from the multicast routing table. However, if you do so, the entries might reappear in the routing table if they are rediscovered.

If you specify an *, the router clears all IP multicast forwarding entries. If you specify the IPv4 address of a multicast group, the router clears all multicast forwarding entries for that group. If you specify the IPv4 address of a multicast group and the IPv4 address of a multicast source, the router clears the multicast forwarding entry that matches that group and source.

To delete IP multicast forwarding entries issue the **clear ip mroute** command in Privileged Exec mode.

```
host1:boston#clear ip mroute *
```

Related Documentation

- Defining Permanent IP Multicast Forwarding Entries on page 8
- Monitoring Multicast Forwarding Entries on page 37
- Monitoring Active Multicast Routes on page 40
- Monitoring Multicast Entries in a Source or Group on page 43
- Monitoring Multicast Routes When OIF Mapping Is Configured on page 43
- Monitoring Multicast Statistics on page 46
- Monitoring Summary Information of Multicast Routes on page 49
- clear ip mroute

Multicast Router Information Support

When you enable multicast routing on a virtual router, the router acts as a multicast router information (mrinfo) server. This feature enables the router to respond to mrinfo requests from other network hosts. Specifically, E Series virtual routers respond to DVMRP ask neighbors and DVMRP ask neighbors2 requests.

Each virtual router responds to mrinfo requests with a list of multicast interfaces and their IP addresses. If appropriate, the virtual router also supplies the following information for each interface:

- Current functional status of the interface (for example, if the interface is down).
- Information as to whether the interface is disabled and the reason for the interface being disabled—either because IP is not configured on the interface or because the interface has been disabled through the software.
- Whether the interface is performing the IGMP queries for this subnet.
- Information about PIM neighbors:

  If PIM is configured on the interface, the virtual router supplies a list of the interface's PIM neighbors and indicates which neighbors are leaf neighbors.
- Information about DVMRP and GRE tunnels:

  If the interface is an endpoint of a tunnel, the virtual router specifies the IP address of the endpoint of the tunnel.

Related Documentation

- IPv4 Multicast Overview on page 4
- Monitoring Summary Information of Multicast Routes on page 49

BGP Multicasting

BGP multicasting (MBGP) is an extension of the BGP unicast routing protocol. Many of the functions available for BGP unicasting are also available for MBGP.

The MBGP extensions specify that BGP can exchange information within different types of *address families*. The address families available are unicast IPv4, multicast IPv4, and VPN-IPv4. When you enable BGP, the router employs unicast IPv4 addresses by default.

We recommend you be thoroughly familiar with BGP before configuring MBGP. See *Configuring BGP Routing* in the *JunosE BGP and MPLS Configuration Guide* for detailed information about BGP and MBGP.

This topic discusses the following:

- Investigating Multicast Routes on page 33

Investigating Multicast Routes

You can use the **mtrace** command to trace the path that multicast packets take from a source to a destination through a multicast group address. This command is similar to the **traceroute** command for investigating unicast routes.

```
host1#mtrace 100.4.4.4 40.1.1.1 232.1.1.1
Tracing multicast route from 100.4.4.4 to 40.1.1.1 for group 232.1.1.1 using response address
  10.6.129.56
(Press ^c to stop.)
Received mtrace response packet of length 88
1. 40.1.1.1 Protocol: PIM(3) FwdingCode: RPF iif(9)
2. 21.2.2.2 Protocol: PIM(3) FwdingCode: Reached RP(8)
```

Related Documentation

- mtrace

CHAPTER 2

Monitoring IPv4 Multicast

The following topics describe how to monitor IP multicast configuration on the E Series router:

Displaying Available Routes for Reverse-Path Forwarding

Purpose Display all available routes, only the routes to a particular destination, or routes associated with a specific unicast protocol that the router can use for RPF. You can specify the IP address and the network mask to view routes to a particular destination. You can also specify a unicast routing protocol to view routes associated with that protocol.

Action To display all routes that the router can use for RPF:

```
host1#show ip rpf-route
Protocol/Route type codes:
  I1- ISIS level 1, I2- ISIS level2,
  I- route type intra, IA- route type inter, E- route type external,
  i- metric type internal, e- metric type external,
  O- OSPF, E1- external type 1, E2- external type2,
  N1- NSSA external type1, N2- NSSA external type2
  L- MPLS label, V- VR/VRF, *- indirect next-hop

  Prefix/Length     Type    Next Hop       Dist/Met    Intf
  -------------     ----    --------       --------    ------
  10.10.0.112/32     Static  192.168.1.1     1/1         fastEthernet0/0
```

```
  10.1.1.0/24      Connect 10.1.1.1       0/1        atm3/0.100
  25.25.25.25/32   Connect 25.25.25.25    0/1        loopback0
```

To display the best static routes added by network management to the routing table:

```
host1#show ip rpf-route static
Protocol/Route type codes:
  I1- ISIS level 1, I2- ISIS level2,
  I- route type intra, IA- route type inter, E- route type external,
  i- metric type internal, e- metric type external,
  O- OSPF, E1- external type 1, E2- external type2,
  N1- NSSA external type1, N2- NSSA external type2
  L- MPLS label, V- VR/VRF, *- indirect next-hop

 Prefix/Length  Type   Next Hop   Dist/Met  Intf
 -------------  ----   --------   --------  --------------
10.10.0.112/32  Static 192.168.1.1  1/1     fastEthernet0/0
```

Meaning Table 5 on page 36 lists the output fields of the **show ip rpf-route** command.

Table 5: show ip rpf-route Output Fields

Field Name	Field Description
Prefix	Value of the logical AND of the IP address of the destination network and the subnet address
Length	Length of the subnet mask in bits
Type	Protocol type for the interface: • Connect—Subnet directly connected to the interface • Static—Static route • *protocol-name*—Route learned through the named protocol
Next Hop	IP address of the next hop for this route
Dist	Distance configured for this route
Met	Learned or configured cost associated with this route
Intf	Type of interface and interface specifier for the next hop. For details about interface types and specifiers, see Interface Types and Specifiers in the *JunosE Command Reference Guide*.

Related Documentation

- Defining Static Routes for Reverse-Path Forwarding on page 7
- Specifying Unicast Routes for RPF on page 8
- show ip rpf-route

Monitoring Multicast Forwarding Entries

Purpose Display information about all or specified multicast routes. You can specify a multicast group IP address or both a multicast group IP address and a multicast source IP address to display information about particular multicast forwarding entries.

Action To display all multicast forwarding entries while bandwidth rate is constant:

```
host1#show ip mroute
IP Multicast Routing Table

(S, G) uptime d h:m:s
  [Data rate: Kbps]  [SPT Threshold: Kbps]  [Threshold: Kbps]
  [Admission bandwidth: bps]
  [QoS bandwidth: bps]
  RPF route: addr/mask, incoming interface
             neighbor address, owner route-owner
  Incoming interface list:
      Interface (addr/mask), State/Owner [(RPF IIF)]
  Outgoing interface list:
      Interface (addr/mask), State/Owner, Uptime/Expires

(10.0.10.1, 225.1.1.1) uptime 0 00:10:31
  Data rate: 2132 Kbps, Threshold 500 Kbps
  Admission bandwidth: 2000000 bps
  RPF route: 10.0.10.0/24, incoming interface atm5/3.1010
             neighbor 10.0.10.8, owner Local
  Incoming interface list:
      atm5/3.1010 (10.0.10.8/24), Accept/Pim (RPF IIF)
  Outgoing interface list:
      atm5/1.108 (108.0.8.5/8), Forward/Pim, 0 00:02:52/never
      atm5/1.109 (107.0.8.4/8), Forward/Pim, 0 00:10:07/never

(1.1.1.1, 225.1.1.1) uptime 0 00:00:34, never expires
RPF route: 1.0.0.0/8, incoming interface ATM5/1.200
             neighbor 2.2.2.2, owner Netmgmt
  Incoming interface list:
      ATM5/1.200 (2.1.1.1/8), Accept/Igmp (RPF IIF)
  Outgoing interface list:
      ATM5/1.300 (3.1.1.1/8), Forward/Igmp, 0 00:00:34/never

Counts:     2 (S, G) entries
            0 (*, G) entries
```

To display all multicasting entries while adaptive bandwidths enabled:

```
Host1#show ip mroute
                  IP Multicast Routing Table

(S, G) uptime d h:m:s[, expires d h:m:s]
  [Admission bandwidth: bps]
  [QoS bandwidth: bps]
  RPF route: addr/mask, incoming interface
             neighbor address, owner route-owner
  Incoming interface list:
      Interface (addr/mask), State/Owner [(RPF IIF)]
  Outgoing interface list:
      Interface (addr/mask), State/Owner, Uptime/Expires

 (10.0.1.9, 225.1.1.1) uptime 0 00:00:23
  Admission bandwidth: 1998000 bps (adaptive)
  QoS bandwidth: 1998000 bps (adaptive)
```



```
  RPF route: 10.0.0.0/8, incoming interface ATM2/1.200
             neighbor 21.1.1.1, owner Netmgmt
  Incoming interface list:
      ATM2/1.200 (21.2.2.2/8), Accept/Pim (RPF IIF)
  Outgoing interface list:
      ATM2/1.300 (31.2.2.2/8), Blocked (port-adm-limit)/Pim, 0 00:00:23/never

Counts: 1 (S, G) entries
        0 (*, G) entries
```

To display all multicast forwarding entries when the bandwidth limit of all the join interfaces exceeds the configured connection admission control limits (CAC):

```
host1#show ip mroute
                  IP Multicast Routing Table

(S, G) uptime d h:m:s[, expires d h:m:s]
  [Admission bandwidth: bps]
  [QoS bandwidth: bps]
  RPF route: addr/mask, incoming interface
             neighbor address, owner route-owner
  Incoming interface list:
      Interface (addr/mask), State/Owner [(RPF IIF)]
  Outgoing interface list:
      Interface (addr/mask), State/Owner, Uptime/Expires

 (10.0.1.9, 225.1.1.1) uptime 0 00:00:23
  Admission bandwidth: 1998724 bps (adaptive)
  QoS bandwidth: 1998724 bps (adaptive)
  RPF route: 10.0.0.0/8, incoming interface ATM2/1.200
             neighbor 21.1.1.1, owner Netmgmt
  Incoming interface list:
      ATM2/1.200 (21.2.2.2/8), Accept/Pim (RPF IIF)
  Outgoing interface list:
     ATM2/1.300 (31.2.2.2/8), Blocked (join-intf-adm-limit)/IGMP, 0 00:00:23/never

Counts: 1 (S, G) entries
        0 (*, G) entries
```

Meaning Table 6 on page 38 lists the output fields of the **show ip mroute** command.

Table 6: show ip mroute Output Fields

Field Name	Field Description
(S, G)	IP addresses of the multicast source and the multicast group
Uptime	Length of time that the (S,G) pair has been active, in *days hours:minutes:seconds* format
Expires	Length of time that the (S,G) pair can be active, in *days hours:minutes:seconds* format or *never*
Data Rate	Flow rate for the threshold entry, in Kbps
SPT Threshold	SPT threshold value for the entry, in Kbps
Threshold	Threshold value for the entry, in Kbps

Table 6: show ip mroute Output Fields *(continued)*

Field Name	Field Description
Admission bandwidth	Admission bandwidth per mroute, in bps
QoS bandwidth	QoS bandwidth per mroute, in bps
RPF Route	IP address and subnetwork mask of the RPF route
incoming interface	Type and specifier of the incoming interface for the RPF route
neighbor address	IP address of the neighbor
State/Owner	Owner of the route: • Local—route belonging to the local interface • Static—Static route • Other protocols—Route established by a protocol such as RIP or OSPF
Incoming interface list	List of incoming interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Accept or Discard • Multicast protocol that owns the interface
Outgoing interface list	List of outgoing interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Forward or Blocked (intf-adm-limit, join-intf-adm-limit, port-adm-limit, port-limit, port-priority-adm-limit) • Protocol running on the interface: PIM, DVMRP, or IGMP • Amount of time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Length of time that the interface can remain active in this multicast forwarding entry, in *days hours:minutes:seconds* format or *never*
Counts	Numbers of types of source group mappings: • (S,G)—Number of (S,G) entries • (*,G)—Number of (*,G) entries

Related Documentation

- Defining Static Routes for Reverse-Path Forwarding on page 7

- Specifying Unicast Routes for RPF on page 8
- Defining Permanent IP Multicast Forwarding Entries on page 8
- Deleting Multicast Forwarding Entries on page 32
- Defining a Multicast Bandwidth Map on page 8
- Blocking Mroutes on page 25
- Creating Mroute Port Limits on page 29

- show ip mroute

Monitoring Active Multicast Routes

Purpose Display active multicast routes.

- You can specify a multicast group IP address or both a multicast group IP address and a multicast source IP address to display information about particular active multicast forwarding entries.
- You can specify the bandwidth threshold to display the active multicast routes with admission bandwidth greater than the specified bandwidth threshold. The default bandwidth threshold is 4000 bps.
- You can use the **summary** option to see a summary rather than a detailed description.
- You can use the **count** option to display the number of active multicast forwarding entries.
- You can use the **oif-detail** option to display the details of the join interfaces corresponding to the mapped interface when oif-mapping is configured.
- You can use the **statistics** option to display statistics for packets received through all active multicast forwarding entries that the router has added to the multicast routing table and established on the appropriate line modules.

Action To display the active multicast routes with bandwidth above 10000 bps:

```
host1#show ip mroute active 10000
              Active IP Multicast Routes >=10000 bps

(S, G) uptime d h:m:s[, expires d h:m:s]
  [Admission bandwidth: bps]
  [QoS bandwidth: bps]
  RPF route: addr/mask, incoming interface
             neighbor address, owner route-owner
  Incoming interface list:
      Interface (addr/mask), State/Owner [(RPF IIF)]
  Outgoing interface list:
      Interface (addr/mask), State/Owner, Uptime/Expires

(52.0.0.1, 232.0.0.1) uptime 0 00:01:07
  Admission bandwidth: 47000 bps (adaptive)
  QoS bandwidth: 47000 bps (adaptive)
  RPF route: 52.0.0.0/24, incoming interface ATM2/1.17
      neighbor 17.0.0.2, owner NetmgmtRpf
  Incoming interface list:
```



```
     ATM2/1.17 (17.0.0.2/24), Accept/Igmp (RPF IIF)
  Outgoing interface list:
     NULL

Counts: 1 (S, G) entries
        0 (*, G) entries
```

To display the summary of active multicast routes:

```
host1#show ip mroute summary active
              Active IP Multicast Routes >=4000 bps

Group Address   Source Address  RPF route          RPF Iif         #Oifs
--------------- --------------- ------------------ --------------- -----
232.0.0.1       51.0.0.1        51.0.0.0/24        ATM3/1.17       0
232.0.0.2       51.0.0.1        51.0.0.0/24        ATM3/1.17       0
232.0.0.3       51.0.0.1        51.0.0.0/24        ATM3/1.17       0
Counts: 3 (S, G) entries
        0 (*, G) entries
```

Meaning Table 7 on page 41 lists the **show ip mroute active** and **show ip mroute summary active** commands output fields.

Table 7: show ip mroute active and show ip mroute summary active Output Fields

Field Name	Field Description
(S,G)	IP addresses of the multicast source and the multicast group
Uptime	Length of time that the (S,G) pair has been active, in *days hours:minutes:seconds* format
Expires	Length of time for which the (S,G) pair will be active, in *days hours:minutes:seconds* format
Admission bandwidth	Admission bandwidth (in bps)
QoS bandwidth	QoS bandwidth (in bps)
RPF Route	IP address and subnetwork mask of the RPF route
Incoming interface	Type and specifier of the incoming interface for the RPF route
neighbor address	IP address of the neighbor
owner	Owner of the route: • Local—route belonging to the local interface • Static—Static route • Other protocols—Route established by a protocol such as RIP or OSPF

Table 7: show ip mroute active and show ip mroute summary active Output Fields *(continued)*

Field Name	Field Description
Incoming interface list	List of incoming interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Accept or Discard • Multicast protocol that owns the interface
Outgoing interface list	List of outgoing interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Forward or Blocked (intf-adm-limit, join-intf-adm-limit, port-adm-limit, port-limit, port-priority-adm-limit) • Protocol running on the interface: PIM, DVMRP, or IGMP • Amount of time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Length of time that the interface can remain active in this multicast forwarding entry, in *days hours:minutes:seconds* format or *never*
Counts	Numbers of types of source group mappings: • (S,G)—Number of (S,G) entries • (*,G)—Number of (*,G) entries
Group Address	IP address of the multicast group
Source Address	IP address of the multicast source
RPF Iif	Type and identifier for the incoming interface for the RPF route
#Oifs	Number of outgoing interfaces

Related Documentation

- Defining Static Routes for Reverse-Path Forwarding on page 7
- Specifying Unicast Routes for RPF on page 8
- Defining Permanent IP Multicast Forwarding Entries on page 8
- Deleting Multicast Forwarding Entries on page 32
- Defining a Multicast Bandwidth Map on page 8
- Blocking Mroutes on page 25

- Creating Mroute Port Limits on page 29
- show ip mroute

Monitoring Multicast Entries in a Source or Group

Purpose Display information about the number of groups and sources.

- You can specify a multicast group IP address or both a multicast group IP address and a multicast source IP address to display information about a particular multicast forwarding entry.
- You can use the **active** option to display information for the active multicast routes.
- You can specify the bandwidth threshold along with the **active** option to display information for the active multicast routes with admission bandwidth greater than the specified bandwidth threshold. The default bandwidth threshold is 4000 bps.

Action To display the number of groups and sources:

```
host1#show ip mroute count
                  IP Multicast Routing Table

Counts:      2 (S, G) entries
             0 (*, G) entries
```

Meaning Table 8 on page 43 lists the **show ip mroute count** command output fields.

Table 8: show ip mroute count Output Fields

Field Name	Field Description
Counts	Number of types of source group mappings: • (S,G)—Number of (S,G) entries • (*,G)—Number of (*,G) entries

Related Documentation

- Defining Permanent IP Multicast Forwarding Entries on page 8
- Deleting Multicast Forwarding Entries on page 32
- Blocking Mroutes on page 25
- show ip mroute

Monitoring Multicast Routes When OIF Mapping Is Configured

Purpose Display details of the join interfaces corresponding to the mapped interfaces when OIF mapping is configured.

- You can specify a multicast group IP address or both a multicast group IP address and a multicast source IP address to display information about a particular multicast forwarding entry.

- You can use the **active** option to display information for the active multicast routes.
- You can specify the bandwidth threshold along with the **active** option to display information for the active multicast routes with admission bandwidth greater than the specified bandwidth threshold. The default bandwidth threshold is 4000 bps.

Action To display details of the join interfaces:

```
host1#show ip mroute oif-detail
                  IP Multicast Routing Table

(S, G) uptime d h:m:s[, expires d h:m:s]
  [Admission bandwidth: bps]
  [QoS bandwidth: bps]
  RPF route: addr/mask, incoming interface
             neighbor address, owner route-owner
  Incoming interface list:
      Interface (addr/mask), State/Owner [(RPF IIF)]
  Outgoing interface list:
      Interface (addr/mask), State/Owner, Uptime/Expires

 (10.0.1.9, 225.1.1.1) uptime 0 00:00:23
  Admission bandwidth: 1998724 bps (adaptive)
  QoS bandwidth: 1998724 bps (adaptive)
  RPF route: 10.0.0.0/8, incoming interface ATM2/1.200
             neighbor 21.1.1.1, owner Netmgmt
  Incoming interface list:
      ATM2/1.200 (21.2.2.2/8), Accept/Pim (RPF IIF)
  Outgoing interface list:
      ATM2/1.300 (31.2.2.2/8), Forward/IGMP, 0 00:00:23/never
    Join interface list:
      ATM2/1.2 (2.2.2.2/32), Blocked (intf-adm-limit)/IGMP, 0 00:00:23/never
      ATM2/1.3 (3.2.2.2/32), Forward/IGMP, 0 00:00:23/never

Counts: 1 (S, G) entries
        0 (*, G) entries
```

Meaning Table 9 on page 44 lists the **show ip mroute oif-detail** command output fields.

Table 9: show ip mroute oif-detail Output Fields

Field Name	Field Description
(S,G)	IP addresses of the multicast source and the multicast group
Uptime	Length of time that the (S,G) pair has been active, in *days hours:minutes:seconds* format
Expires	Length of time for which the (S,G) pair will be active, in *days hours:minutes:seconds* format
Admission bandwidth	Admission bandwidth (in bps)
QoS bandwidth	QoS bandwidth (in bps)
RPF Route	IP address and subnetwork mask of the RPF route

Table 9: show ip mroute oif-detail Output Fields *(continued)*

Field Name	Field Description
Incoming interface	Type and specifier of the incoming interface for the RPF route
neighbor address	IP address of the neighbor
owner	Owner of the route: • Local—Route belonging to the local interface • Static—Static route • Other protocols—Route established by a protocol such as RIP or OSPF
Incoming interface list	List of incoming interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Accept or Discard • Multicast protocol that owns the interface
Outgoing interface list	List of outgoing interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Forward or Blocked (intf-adm-limit, join-intf-adm-limit, port-adm-limit, port-limit, port-priority-adm-limit) • Protocol running on the interface: PIM, DVMRP, or IGMP • Amount of time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Length of time that the interface can remain active in this multicast forwarding entry, in *days hours:minutes:seconds* format or *never*
Join interface list	List of join interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Forward or Blocked (intf-adm-limit, join-intf-adm-limit, port-adm-limit, port-limit, port-priority-adm-limit) • Protocol running on the interface: PIM, DVMRP, or IGMP • Amount of time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Length of time that the interface can remain active in this multicast forwarding entry, in *days hours:minutes:seconds* format or *never*

Table 9: show ip mroute oif-detail Output Fields *(continued)*

Field Name	Field Description
Counts	Numbers of types of source group mappings: • (S,G)—Number of (S,G) entries • (*,G)—Number of (*,G) entries

Related Documentation

- Defining Static Routes for Reverse-Path Forwarding on page 7
- Defining Permanent IP Multicast Forwarding Entries on page 8
- Deleting Multicast Forwarding Entries on page 32
- Defining a Multicast Bandwidth Map on page 8
- Configuring Hardware Multicast Packet Replication With OIF-Mapping on page 25
- Blocking Mroutes on page 25
- Creating Mroute Port Limits on page 29

- show ip mroute

Monitoring Multicast Statistics

Purpose Display statistics for packets received through multicast routes that the router has added to the multicast routing table and established on the appropriate line modules.

- You can specify a multicast group IP address or both a multicast group IP address and a multicast source IP address to display information about a particular multicast forwarding entry.
- You can use the **active** option to display information for the active multicast routes.
- You can specify the bandwidth threshold along with the **active** option to display information for the active multicast routes with admission bandwidth greater than the specified bandwidth threshold. The default bandwidth threshold is 4000 bps.

Action To display statistics of the multicast routes added to the multicast routing table:

```
host1#show ip mroute statistics
IP Multicast Routing Table

(S, G) uptime d h:m:s[, expires d h:m:s]
  [Admission bandwidth: bps]
  [QoS bandwidth: bps]
  RPF route: addr/mask, incoming interface
             neighbor address, owner route-owner
  Incoming interface list:
      Interface (addr/mask), State/Owner [(RPF IIF)]
  Outgoing interface list:
      Interface (addr/mask), State/Owner, Uptime/Expires
(10.0.1.9, 225.1.1.1) uptime 0 00:00:23
  Admission bandwidth: 2000000 bps
  QoS bandwidth: 2000000 bps
```



```
  RPF route: 10.0.0.0/8, incoming interface ATM2/1.200
             neighbor 21.1.1.1, owner Netmgmt
  Incoming interface list:
      ATM2/1.200 (21.2.2.2/8), Accept/Pim (RPF IIF)
  Outgoing interface list:
      ATM2/1.300 (31.2.2.2/8), Blocked (port-adm-limit)/Pim, 0 00:00:23/never
  Statistics:
      Received   : 23 pkts, 1472 bytes
      Forwarded  : 0 pkts, 0 bytes
      Rcvd on OIF: 0 pkts

Counts: 1 (S, G) entries
        0 (*, G) entries
```

Meaning Table 10 on page 47 lists the **show ip mroute statistics** command output fields.

Table 10: show ip mroute statistics Output Fields

Field Name	Field Description
(S,G)	IP addresses of the multicast source and the multicast group
Uptime	Length of time that the (S,G) pair has been active, in *days hours:minutes:seconds* format
Expires	Length of time for which the (S,G) pair will be active, in *days hours:minutes:seconds* format
Admission bandwidth	Admission bandwidth (in bps)
QoS bandwidth	QoS bandwidth (in bps)
RPF Route	IP address and subnetwork mask of the RPF route
Incoming interface	Type and specifier of the incoming interface for the RPF route
neighbor address	IP address of the neighbor
owner	Owner of the route: • Local—Route belonging to the local interface • Static—Static route • Other protocols—Route established by a protocol such as RIP or OSPF
Incoming interface list	List of incoming interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Accept or Discard • Multicast protocol that owns the interface

Table 10: show ip mroute statistics Output Fields *(continued)*

Field Name	Field Description
Outgoing interface list	List of outgoing interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Forward or Blocked (intf-adm-limit, join-intf-adm-limit, port-adm-limit, port-limit, port-priority-adm-limit) • Protocol running on the interface: PIM, DVMRP, or IGMP • Amount of time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Length of time that the interface can remain active in this multicast forwarding entry, in *days hours:minutes:seconds* format or *never*
Statistics	• Received—Number of packets and bytes that the virtual router received for this multicast route • Forwarded—Number of packets and statistics that the virtual router has forwarded for this multicast route • Rcvd on OIF—Number of packets and statistics that the virtual router has received on the OIF for this multicast route NOTE: The output shows statistics after the virtual router has added the multicast route to the multicast routing table and established the route on the appropriate line module. Statistics for interactions before the route is established on the line module are not displayed.
Counts	Numbers of types of (S,G) mappings: • (S,G)—Number of (S,G) entries • (*,G)—Number of (*,G) entries

Related Documentation

- Defining Static Routes for Reverse-Path Forwarding on page 7
- Specifying Unicast Routes for RPF on page 8
- Defining Permanent IP Multicast Forwarding Entries on page 8
- Deleting Multicast Forwarding Entries on page 32
- Defining a Multicast Bandwidth Map on page 8
- Blocking Mroutes on page 25
- Creating Mroute Port Limits on page 29

- show ip mroute

Monitoring Summary Information of Multicast Routes

Purpose Display summary of all or specified multicast routes.

- You can specify a multicast group IP address or both a multicast group IP address and a multicast source IP address to display information about a particular multicast forwarding entry.
- You can use the **active** option to display information for the active multicast routes.
- You can specify the bandwidth threshold along with the **active** option to display information for the active multicast routes with admission bandwidth greater than the specified bandwidth threshold. The default bandwidth threshold is 4000 bps.

Action To display brief information about the multicast routes:

```
host1#show ip mroute summary
                        IP Multicast Routing Table

Group Address    Source Address   RPF route    RPF Iif          #Oifs
---------------  ---------------  ---------    -------------    ----
224.0.1.39       52.1.1.1         51.1.1.1/32  Register IIF     0
224.0.1.40       51.1.1.1         51.1.1.1/32  loopback1        1

Counts:     2 (S, G) entries
            0 (*, G) entries
```

Meaning Table 11 on page 49 lists the **show ip mroute summary** command output fields.

Table 11: show ip mroute summary Output Fields

Field Name	Field Description
Group Address	IP address of the multicast group
Source Address	IP address of the multicast source
RPF Route	IP address and network mask of the RPF route
RPF Iif	Type and identifier for the incoming interface for the RPF route
#Oifs	Number of outgoing interfaces
Counts	Numbers of types of (S,G) pairs: • (S,G)—Number of (S,G) entries • (*,G)—Number of (*,G) entries

Related Documentation

- Defining Static Routes for Reverse-Path Forwarding on page 7
- Specifying Unicast Routes for RPF on page 8
- Defining Permanent IP Multicast Forwarding Entries on page 8

- Deleting Multicast Forwarding Entries on page 32
- Blocking Mroutes on page 25
- Creating Mroute Port Limits on page 29

- show ip mroute

Monitoring Multicast Protocols Enabled on the Router

Purpose Display information about the multicast protocols enabled on the router.

Action To display information about the multicast protocols enabled on the router:

```
host1#show ip multicast protocols
Multicast protocols:

Protocol Pim
    Type: Sparse
    Interfaces: 1 registered, 1 owned
    Registered interfaces:
        ATM2/1.103 (103.0.0.2/24) owner Pim

Protocol Igmp
    Type: Local
    Interfaces: 1000 registered, 1000 owned
    Registered interfaces:
        ATM2/0.131 (13.0.0.1/24) local Igmp owner Igmp
            Admission-bandwidth 2000000/10000000 bps
            QoS Adjust 2000000 bps
        Active <S,G> count   15
        Blocked <S,G> count  10

        ATM2/0.132 (13.0.0.2/24) local Igmp owner Igmp
            Admission-bandwidth 0/10000000 bps
            QoS Adjust 0 bps
        Active <S,G> count   25
        Blocked <S,G> count  10

        ATM2/0.133 (13.0.0.3/24) local Igmp owner Igmp
            Admission-bandwidth 8000000/10000000 bps
            QoS Adjust 0 bps
...
Count: 2  protocols
```

Meaning Table 12 on page 50 lists the **show ip multicast protocols** command output fields.

Table 12: show ip multicast protocols Output Fields

Field Name	Field Description
Multicast Protocols	Multicast protocols on this router
Protocol	Name of the multicast protocol

Table 12: show ip multicast protocols Output Fields *(continued)*

Field Name	Field Description
Type	Mode of the multicast protocol: • For DVMRP—Dense • For PIM—Sparse, Dense, or Sparse-Dense • For IGMP—Local
Interfaces	• registered—Number of interfaces on which the protocol is configured • owned—Number of interfaces that a protocol owns. If you configure only IGMP on an interface, IGMP owns the interface. However, if you configure IGMP and either PIM or DVMRP on the same interface, PIM or DVMRP owns the interface.
Registered interfaces	Information about interfaces on which the protocol is configured: • Types and identifiers of interfaces. For details about interface types and specifiers, see Interface Types and Specifiers in the *JunosE Command Reference Guide*. • Protocols configured on the interface and the protocol that owns the interface. If you configure only IGMP on an interface, IGMP owns the interface. However, if you configure IGMP and PIM or DVMRP on the same interface, PIM or DVMRP owns the interface. • Admission-bandwidth—Actual admission bandwidth or configured admission bandwidth (in bps) • QoS Adjust—Bandwidth of QoS adjustment, in bps • Count—Number of multicast protocols on the virtual router • Active <S,G> count—Number of active S,G data streams on the interface • Blocked <S,G> count—Number of blocked S,G data streams on the interface

Related Documentation

- Defining a Multicast Bandwidth Map on page 8
- Blocking Mroutes on page 25
- Enabling Interface Admission Bandwidth Limitation on page 26
- Enabling Port-Level Admission Bandwidth Control on page 29

- show ip multicast protocols

Monitoring Summary Information of Multicast Protocols Enabled on the Router

Purpose Display a summary of information about the multicast protocols enabled on the router.

Action To display a summary of information about the multicast protocols enabled on the router:

```
host1#show ip multicast protocols brief
Protocol  Registered  Owned       Type
          Interfaces  Interfaces
--------- ----------  ----------  --------------------
Pim       2           2           Sparse Dense
Igmp      1           0           Local

Count: 2 protocols
```

Meaning Table 13 on page 52 lists the **show ip multicast protocols brief** command output fields.

Table 13: show ip multicast protocols brief Output Fields

Field Name	Field Description
Protocol	Name of the multicast protocol
Registered Interfaces	Number of interfaces on which the protocol is configured
Owned Interfaces	Number of interfaces that a protocol owns. If you configure only IGMP on an interface, IGMP owns the interface. However, if you configure IGMP and either PIM or DVMRP on the same interface, PIM or DVMRP owns the interface.
Type	Mode of the multicast protocol: • For DVMRP—Dense • For PIM—Sparse, dense, or sparse-dense • For IGMP—Local
Count	Number of multicast protocols on the virtual router

Related Documentation

- show ip multicast protocols

Monitoring the IP Multicast Status on a Virtual Router

Purpose Display information about the status of IP multicast on the virtual router.

Action To display information about the status of IP multicast on the virtual router:

```
host1#show ip multicast routing
Multicast forwarding is enabled on this router
Multicast graceful restart is complete (timer 0 seconds) on this router
Multicast cache-miss processing is enabled on this router
```

Related Documentation

- Enabling IP Multicast on page 7
- Enabling and Disabling RPF Checks on page 8
- Defining Permanent IP Multicast Forwarding Entries on page 8

- Activating Multicast QoS Adjustment Functions on page 16
- show ip multicast routing

Monitoring Multicast Routes on Virtual Router Ports

Purpose Display information for multicast routes on a port across all virtual routers.

NOTE: This command displays information for mroutes on a port across all virtual routers.

Action To display the multicast route port outgoing interface, limits, counts, bandwidth settings, and bandwidth accepted:

host1#**show mroute port count**

```
BW     Priority
Port    Limit   Count    bps     BW bps    Hysteresis   Admitted
-----   -----   -----   -----   --------   ----------   --------
1/1/0   None    1       None    None       85           0
1/1/1   None    2       15000   10000      85           2000
```

Meaning Table 14 on page 53 lists the output fields of the **show mroute port count** command.

Table 14: show mroute port count Output Fields

Field Name	Field Description
Port	Slot or port value on the router
Limit	None (reserved for future functionality)
Count	Number of multicast route outgoing interfaces on the specified port
BW bps	Bandwidth limit (in bits per second)
Priority BW bps	Priority bandwidth limit (in bits per second)
Admitted	Bandwidth admitted on the port (in bits per second)

Related Documentation

- Defining a Multicast Bandwidth Map on page 8
- Creating Mroute Port Limits on page 29
- Enabling Port-Level Admission Bandwidth Control on page 29
- show mroute port count

CHAPTER 3

Configuring IGMP and IGMP Proxy

IP hosts use Internet Group Management Protocol (IGMP) in IPv4 to report their multicast group memberships to neighboring routers. Similarly, multicast routers, such as an E Series router, use IGMP to discover which of their hosts belong to multicast groups.

This chapter describes how to configure IGMP for IP multicast on an E Series router; it contains the following sections:

IGMP Overview

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

This topic discusses the following:

- Group Membership Queries on page 56
- Group Membership Reports on page 57
- Leave Group Membership Messages on page 57

Group Membership Queries

A multicast router can be a querier or a nonquerier. Only one querier is on a network at any time. Multicast routers monitor queries from other multicast routers to determine the status of the querier. If the querier detects a query from a router with a lower IP address, it relinquishes its role to that router.

IGMPv1 and IGMPv2 mode interfaces send two types of group membership queries to hosts on the network:

- General queries to the all-hosts group address (224.0.0.1)
- Specific queries to the appropriate multicast group address

IGMPv3 mode interfaces send the following types of queries to IGMPv3 hosts:

- General queries
- Group-specific queries
- Source-specific queries

The purpose of a group membership query is to discover the multicast groups to which a host belongs.

IGMPv2 and IGMPv3 group membership queries have a Max Response Time field. This response time is the maximum amount of time that a host can take to reply to a query.

Group Membership Reports

When a host receives a group membership query, it identifies the groups associated with the query and determines to which groups the query belongs. The host then sets a timer, with a value less than the Max Response Time field in the query, for each group to which it belongs.

When the timer expires, the host sends a group membership report to the group address. When a multicast router receives a report, it adds the group to the membership list for the network and sets a timer to the *group membership interval*. The router calculates the group membership interval using the following formula of configurable IGMP values:

(query interval x robustness value) + query maximum response time

If this timer interval expires before the router receives another group membership report, the router determines that the group has no members left on the network.

IGMPv3 supports an extended report format you can use to report multiple groups and source lists in a single report.

Leave Group Membership Messages

When a host leaves a group, it sends a leave group membership message to multicast routers on the network. A host generally addresses leave group membership messages to the all-routers group address (224.0.0.2).

Related Documentation

- IGMP Platform Considerations on page 57
- IGMP References on page 58
- Enabling IGMP on an Interface on page 60

IGMP Platform Considerations

For information about modules that support IGMP on the ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support IGMP.

For information about modules that support IGMP on the E120 and E320 Broadband Services Routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support IGMP.

Related Documentation

- IGMP Overview on page 56
- IGMP References on page 58

IGMP References

For more information about IGMP, see the following resources:

- IGMP-based Multicast Forwarding (“ IGMP Proxying”)—draft-ietf-magma-igmp-proxy-00.txt (May 2002 expiration)
- RFC 2236—Internet Group Management Protocol, Version 2 (November 1997)
- RFC 2933—Internet Group Management Protocol MIB (October 2000)
- RFC 3292—General Switch Management Protocol (GSMP) V3 (June 2002)
- RFC 3376—Internet Group Management Protocol (October 2002)
- GSMP extensions for layer2 control (L2C) Topology Discovery and Line Configuration—draft-wadhwa-gsmp-l2control-configuration-00.txt (July 2006 expiration)

Related Documentation

- IGMP Overview on page 56
- IGMP Platform Considerations on page 57

Static and Dynamic IGMP Interfaces

The router supports *static* and *dynamic* IGMP interfaces. Unlike static interfaces, dynamic interfaces are not restored when you reboot the router. For some protocols, dynamic layers can build on static layers in an interface; however, in a dynamic IGMP interface, all the layers are dynamic. See Figure 6 on page 59 for examples of static and dynamic IGMP interfaces.

Figure 6: Static and Dynamic IGMP Interfaces

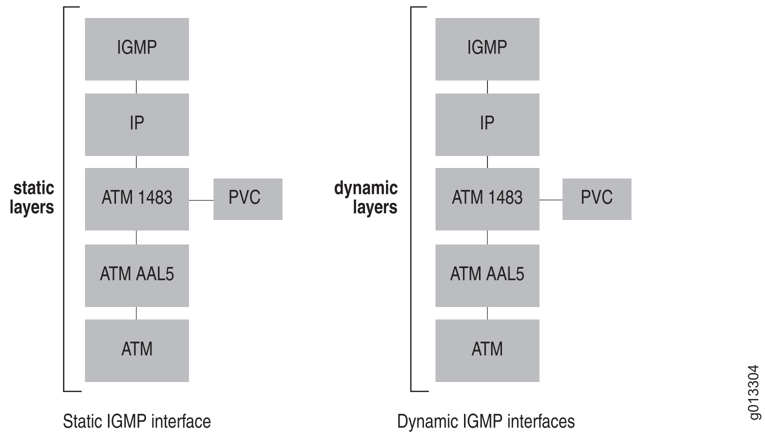

You configure static IGMP interfaces by using software such as the CLI or an SNMP application; you configure dynamic IGMP interfaces by using a profile. A profile constitutes a set of attributes for an interface; a profile for dynamic IGMP interfaces contains attributes for configuring all the layers in the interface.

You define a profile by using the same CLI commands that you use to configure a static IGMP interface; however, the mode in which you use the commands differs. Use the commands in Interface Configuration mode to configure a static IGMP interface and in Profile Configuration mode to define a profile.

When you have defined a profile, you can apply it to an interface or a group of interfaces. Profiles provide an efficient method of creating and managing large numbers of dynamic interfaces. For detailed information about creating and assigning profiles, see *Configuring Dynamic Interfaces* in the *JunosE Link Layer Configuration Guide*. When you create a profile for dynamic IGMP interfaces, specify attributes for configuring all layers in the interface.

You use the following IGMP commands to configure a static IGMP interface. You also use these commands to define the attributes for the IGMP layer when you create a profile for dynamic IGMP interfaces:

Table 15: IGMP Commands

ip igmp	ip igmp query-max-response-time
ip igmp access-group	ip igmp robustness
ip igmp access-source-group	ip igmp ssm-map enable
ip igmp apply-oif-map	ip igmp ssm-map static
ip igmp explicit-tracking	ip igmp query-interval
ip igmp group limit	ip igmp static-exclude
ip igmp immediate-leave	ip igmp static-group

Table 15: IGMP Commands *(continued)*

ip igmp last-member-query-interval	ip igmp static-include
ip igmp promiscuous	ip igmp version
ip igmp querier	multicast group port limit
ip igmp querier-timeout	

The following sections describe the tasks associated with these and other **ip igmp** commands.

You can also use various IGMP-specific RADIUS attributes in RADIUS Access-Accept messages as an alternative method of configuring certain values. See *Configuring RADIUS Attributes* in the *JunosE Broadband Access Configuration Guide* for additional information.

Related Documentation

- Dynamic Interfaces Overview

Before You Begin Configuring IGMP

You can configure IGMP on IPv4 multicast interfaces.

For information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- Configuring IPv4 Multicast Attributes on page 7
- Configuring IPv6 Multicast Attributes on page 165

Enabling IGMP on an Interface

You must start IGMP on each interface that you want to use the protocol. You can configure IGMP and either PIM or DVMRP on the same interface. If you configure only IGMP on an interface, IGMP owns that interface. If you configure IGMP and either PIM or DVMRP on an interface, PIM or DVMRP owns the interface.

By enabling IGMP, the router processes incoming multicast packets and creates an entry in the multicast routing table. If neither PIM nor DVMRP own the interface (for example, when only IGMP is configured), then the packets are locally routed to other interfaces on the router. PIM or DVMRP must be configured on the interface for packets to be sent to other routers.

For networks that use only IGMPv1, you can configure an interface to operate in IGMPv1 mode. However, IGMPv2 and IGMPv3 interfaces support IGMPv1 hosts. In an IGMPv1 network, you must configure one interface to act as a querier. In an IGMPv2 or IGMPv3 network, the querier is the router with the lowest IP address.

To start IGMP, complete the following steps:

1. Enable IGMP on the interface (IGMPv2 is the default version).

   host1:boston(config-if)#**ip igmp**

2. (IGMPv1 or IGMPv3) Specify the IGMP version for the interface.

   host1:boston(config-if)#**ip igmp version 1**

3. (IGMPv1 only) Specify that the interface act as the querier for the network.

   host1:boston(config-if)#**ip igmp querier**

Related Documentation

- Configuring IGMP Settings for an Interface on page 61
- Configuring Multicast Groups for IGMP on page 62
- ip igmp

Configuring IGMP Settings for an Interface

When you start IGMP on an interface, it operates with the default settings. You can, however, modify:

- The method that the router uses to remove hosts from multicast groups (IGMPv2 and IGMPv3 interfaces only).

  host1:boston(config-if)#**ip igmp immediate-leave**

- The query time interval for the querier sends group membership messages.

  host1:boston(config-if)#**ip igmp last-member-query-interval 90**

- The time that a non-querier waits for queries from the current querier before sending query messages to assume responsibility of querier.

  host1:boston(config-if)#**ip igmp querier-timeout 200**

- The time that a new querier waits before sending query messages after it assumes responsibility from another querier.

  host1:boston(config-if)#**ip igmp query-interval 100**

- The time that a host can take to reply to a query (maximum response time).

  host1:boston(config-if)#**ip igmp query-max-response-time 120**

- The number of times that the router sends each IGMP message from this interface.

  host1:boston(config-if)#**ip igmp robustness 2**

Related Documentation

- Enabling IGMP on an Interface on page 60
- Configuring Multicast Groups for IGMP on page 62
- ip igmp immediate-leave
- ip igmp last-member-query-interval
- ip igmp querier-timeout
- ip igmp query-interval

- ip igmp query-max-response-time
- ip igmp robustness

Configuring Multicast Groups for IGMP

You can configure multicast groups for IGMP using the following tasks:

- Specifying Multicast Groups on page 62
- Assigning a Multicast Group to an Interface on page 62
- Configuring Group Outgoing Interface Mapping on page 62

Specifying Multicast Groups

You can use a standard-format or extended-format IP access list to specify the multicast groups that a host can join.

1. Configure the host on a subnetwork to join only multicast groups that appear on the specified IP access list.

   host1:boston(config-if)#**ip igmp access-group boston-list**

2. Configure the host on a subnetwork to become a member of only those (S,G) pairs (also known as *channels*) which are included on the specified IP access list.

   host1:boston(config-if)#**ip igmp access-source-group dallas-list**

Assigning a Multicast Group to an Interface

You can assign an interface to send and receive all traffic for a particular multicast group. This feature enables you to control the IGMP traffic and to test the behavior of multicast protocols in the network.

To send and receive all traffic for a multicast group from a specific interface:

- Issue the **ip igmp static-group** command in Interface Configuration mode.

  host1:boston(config-if)#**ip igmp static-group 225.1.2.3**

You can use the **no** version to stop the interface from sending all traffic for the group. The interface sets no timers for this group.

Configuring Group Outgoing Interface Mapping

You can configure an IGMP interface to use a different outgoing interface (OIF) for multicast-data-forwarding by applying an OIF map. When you configure an OIF map on an IGMP interface, the map is applied to all IGMP membership requests that the interface receives.

To configure OIF mapping on an interface:

1. Create an OIF map using the **ip igmp oif-map** command at the global level.

   host1(config)#**ip igmp oif-map OIFMAP atm 3/0.1 232.0.0.1/32 51.0.0.1/32**

host1(config)#ip igmp oif-map OIFMAP atm 3/0.2 232.0.0.1/32 51.0.0.2/32
host1(config)#ip igmp oif-map OIFMAP atm 3/0.3 233.0.0.1/32
host1(config)#ip igmp oif-map OIFMAP atm 3/0.4 233.0.0.0/24 51.0.0.1/32
host1(config)#ip igmp oif-map OIFMAP atm 3/0.5 233.0.0.0/24 51.0.0.2/32
host1(config)#ip igmp oif-map OIFMAP self 0.0.0.0/0 51.0.0.0/24

2. Apply the OIF map to an interface using the **ip igmp apply-oif-map** command.

   host1(config-subif)#ip igmp apply-oif-map OIFMAP

To properly configure an interface used in the OIF map for multicast-data-forwarding capability, you must configure the interface version as passive with the **ip igmp version** command. You can either specify a passive interface as the OIF or specify the OIF as *self* (to use the IGMP interface as the OIF) in the **ip igmp oif-map** command.

Access Node Control Protocol for IGMP

By using ANCP, IGMP is no longer terminated or proxied at the access node. Instead, IGMP passes through the access node transparently. B-RAS terminates both the data PVC and IGMP. After possible user permission verification, B-RAS may instruct the access node, by using GSMP, to establish a multicast branch for the subscriber port.

Access Node Control Protocol (ANCP), also known as Layer 2 Control (L2C) works with a special IGMP session to collect OIF mapping events in a scalable manner. For additional information about configuring ANCP for IGMP, see *Configuring ANCP* in the *JunosE IP Services Configuration Guide*.

Related Documentation

- Configuring ANCP
- Configuring Group Outgoing Interface Mapping on page 62

SSM Mapping

Source-specific multicast (SSM) mapping enables the router to determine one or more source addresses for group G. The mapping effectively translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report, enabling the router to continue as if it had initially received an IGMPv3 report. After the router is joined to these groups, it sends out PIM join messages and continues to enable joining from these groups, as long as it continues to receive IGMPv1 and IGMPv2 membership reports and no change occurs to the SSM mapping for the group.

When you statically configure SSM mapping, the router can discover source addresses from a statically configured table.

The following conditions apply when you configure SSM mapping:

- When an SSM map is configured without any matching access list, SSM mapping is not applied on the incoming (*,G) groups. The PIM SSM range must deny any unacceptable SSM group addresses.

NOTE: An access list must be explicitly configured with the same name as that of the SSM map and group addresses that are to be SSM mapped.

- When you issue the **no ip igmp ssm-map enable** command, the router removes all SSM map (S,G) states and establishes a (*,G) state.
- You can enter multiple **ssm-map static** commands for different access lists. Also, you can enter multiple **ssm-map static** commands for the same access list, as long as the access list uses different source addresses.
- SSM maps do not process statically configured groups.

Related Documentation

- ip igmp ssm-map static

Overview of Limiting the Number of Accepted IGMP Groups

By default, there is no limit on the number of IGMP groups that an IGMP interface can accept. However, you can manage multicast traffic on the router by restricting the number of IGMP groups accepted by:

- A specific port on an I/O module
- A specific IGMP interface

If you set limits for both a port and interfaces on that port, the router uses the lower of the two limits when determining how many IGMP groups an interface can accept. For example, if you set a limit of 10 groups for the port and 15 groups for each interface, only 10 groups can be accepted among the interfaces.

However, if you set a limit for a port and that limit is lower than the number of groups currently accepted by the interfaces on that port, the router does not dissociate the groups from the interfaces. The router enforces the new limit on the port when the number of groups associated with the interfaces falls to that limit. For example, if the interfaces on the port have accepted a total of 15 groups, and you set a limit of 10 groups on the port, the router does not disconnect any of the groups and prevents the interfaces from accepting any more groups. Over time, some groups leave the interfaces and, eventually, a maximum of ten groups remain connected.

Related Documentation

- Configuring Multicast Groups for IGMP on page 62

IGMP Traffic Overview

IGMPv3 extends IGMPv2 functionality with the ability to include or exclude specific multicast traffic sources. That is, with IGMPv3, hosts signal (S,G) pairs to be included or excluded.

For hosts that cannot signal group membership dynamically, you can use the **ip igmp static-include** or **ip igmp static-exclude** command to statically include or exclude multicast traffic, respectively.

IGMPv3 is the industry-designated standard protocol for hosts to signal channel subscriptions in SSM.

Related Documentation

- Understanding PIM Source-Specific Mode
- ip igmp static-include
- ip igmp static-exclude

Explicit Host Tracking

Explicit host tracking enables the router to explicitly track each individual host that is joined to a group or channel on a particular multi-access network.

Explicit host tracking provides the following:

- Minimal leave latency when a host leaves a multicast group or channel. When the router receives a leave message for a group or channel on an interface, it accesses a list of hosts and immediately stops forwarding traffic if the sender is the last host to request traffic for that group or channel. The leave latency is bound only by the packet transmission latencies in the multi-access network and the processing time in the router.
- Ability to change channels quickly in networks where bandwidth is constrained between a multicast-enabled router and hosts.
- Ability to determine what multicast hosts are joined to particular multicast groups or channels, which is useful for accounting purposes.
- Reduction of control message traffic on the network because, when it receives a leave message, the router no longer needs to send out IGMP queries to verify membership. As a result, interested hosts also do not need to respond to these queries with reports.
- Tracking based on the IGMP reports for hosts in both include and exclude modes for every multicast group or channel on an interface.

When the router is configured for explicit host tracking and starts immediate leave using the host information collected, every leave message received for a group or channel is treated as follows:

- The router checks the number of hosts that receive traffic from the group or channel.

- If the host sending the leave message is the only host, it starts immediate leave for that group or channel on that interface. The router removes the interface from the multicast group or channel immediately, without sending out a group or group-source-specific query and waiting for the last member query interval.
- If the host sending the leave message is not the only host receiving traffic for that group or channel, the router removes the host from the list of hosts on that interface, but keeps the interface in the outgoing interface list for the multicast group or channel. No group or group-source-specific queries are sent.

If one or more hosts that support only IGMP V1 are present on a network, the leave latencies for the multicast groups to which those hosts are joined revert to the IGMP V1 leave latency. This affects only the multicast groups to which these legacy hosts are actually joined at any point in time.

You cannot configure explicit host tracking on passive IGMP interfaces or on IGMP V1 interfaces. When you enable IGMP V2 or V3 on an interface, explicit host tracking is not enabled by default.

When you enable explicit host tracking on an interface that has a membership state, the router does not immediately start performing immediate leave. For a maximum of group membership interval seconds, the router only performs host tracking. Any leave messages that the router receives during this period receive normal leave processing. Any leave messages received after this interval has elapsed receive immediate leave processing, when appropriate.

When explicit host tracking has been enabled on an IGMP V3 interface, even if a group has to downgrade to IGMP V2 due to the presence of an IGMP V2 host, explicit host tracking continues for that group. To avoid this, you can use the disable-if-igmp-v2-detected keyword. If you select this option, the router turns off explicit host tracking for the group when IGMP V2 host reports are received for the group on that interface. This option does not have any significance on an interface configured for IGMP V2 and is ignored if provided. Because IGMP V1 does not support leave messages, explicit host tracking is turned off for a group that downgrades to IGMP V1 due to the presence of IGMP V1 hosts.

Explicit host tracking cannot be enabled on an interface that has immediate-leave configured and vice versa. Any attempt to configure immediate-leave on an interface that has explicit host tracking enabled or to configure explicit host tracking on an interface that has immediate-leave enabled is rejected and an error message logged on the screen.

Related Documentation

- Configuring Explicit Host Tracking on page 67

Configuring IGMP Attributes

You can configure IGMP attributes using the following tasks:

- Configuring Explicit Host Tracking on page 67
- Disabling and Removing IGMP on page 67

Configuring Explicit Host Tracking

The following example enables IGMP V3 explicit host tracking on interface 3/0.101 with the default configuration where the router continues to perform explicit host tracking for IGMP V2 groups. To override this default configuration, you must use the **ip igmp explicit-tracking disable-if-igmp-v2-detected** command.

```
user@host#interface 3/0.101
user@host#ip igmp version 3
user@host#ip igmp explicit-tracking
```

Disabling and Removing IGMP

To disable and re-enable IGMP on the VR:

- Disable IGMP on a VR.

  ```
  host1(config)#virtual-router boston
  host1:boston(config)#router igmp
  host1:boston(config-router)#igmp disable
  ```

- Enable IGMP on a VR.

  ```
  host1(config)#virtual-router boston
  host1:boston(config)#router igmp
  ```

You can remove IGMP from the VR and recreate it on the VR.

Related Documentation

- igmp disable
- router igmp
- virtual-router

Example: Accepting IGMP Reports from Remote Subnetworks

This example shows how to accept IGMP reports from a remote subnetwork.

- Requirements on page 67
- Overview on page 68
- Configuration on page 68

Requirements

This example uses the following hardware and software components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

By default, IGMP interfaces accept IGMP reports only from associated subnetworks. You can configure the router to accept IGMP reports from subnetworks that are not associated with its interfaces. The **igmp promiscuous** command in Router Configuration mode specifies whether interfaces on the router can accept IGMP reports from indirectly connected subnets. To override this global setting on a particular interface, use the **ip igmp promiscuous** command in Interface Configuration mode.

Configuration

In the following example, the router is configured to accept IGMP reports from indirectly connected subnets on all interfaces. The interface on port 0 of the line module in slot 4 is then configured to accept IGMP reports only from directly connected subnets. This configuration discusses the following procedure:

- Configuring Router Interface to Accept IGMP Reports from Indirectly Connected Subnets on page 68

Configuring Router Interface to Accept IGMP Reports from Indirectly Connected Subnets

Step-by-Step Procedure

1. Access virtual-router boston and enable IGMP on the virtual-router.

   host1(config)#**virtual-router boston**

2. Enable all IGMP interfaces on the router to accept IGMP reports from hosts on any subnetwork.

   host1:boston(config-router)#**igmp promiscuous**
   host1:boston(config-router)#**exit**

3. Enable the interface to accept IGMP reports only from hosts on subnetworks associated with this interface.

   host1:boston(config)#**interface serial 4/0**
   host1:boston(config-if)#**ip igmp promiscuous off**

Related Documentation

- ip igmp promiscuous
- virtual-router

IGMP Proxy Overview

IGMP proxy enables the router to issue IGMP host messages on behalf of hosts that the router discovered through standard IGMP interfaces. The router acts as a *proxy* for its hosts. E Series routers support IGMP proxy versions 2 and 3.

Figure 7 on page 69 shows a router in an IGMP proxy configuration. You enable IGMP proxy on one interface, which connects to a router closer to the root of the tree. This interface is the *upstream interface*. The router on the upstream interface is running IGMP.

You enable IGMP on the interfaces that connect the router to its hosts that are farther away from the root of the tree. These interfaces are known as *downstream interfaces*.

Figure 7: Upstream and Downstream Interfaces

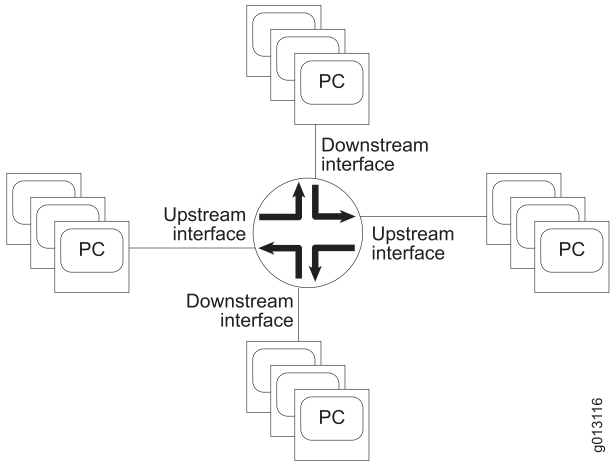

As described in "IGMP Overview" on page 56, earlier in this chapter, hosts interact with the router through the exchange of IGMP messages. Similarly, when you configure IGMP proxy, the router interacts with the router on its upstream interface through the exchange of IGMP messages. However, when acting as the proxy, the router performs the host portion of the IGMP task on the upstream interface, as follows:

- When queried, sends group membership reports to the group.
- When one of its hosts joins a multicast address group to which none of its other hosts belong, sends unsolicited group membership reports to that group.
- When the last of its hosts in a particular multicast group leaves the group, sends an unsolicited leave group membership report to the all-routers group (244.0.0.2).

Related Documentation

- Configuring IGMP Proxy on page 69

Configuring IGMP Proxy

To configure a downstream interface, enable IGMP on that interface. To configure IGMP proxy on the router, complete the following tasks:

1. Enable IP multicast.

   host1(config)#**ip multicast-routing**

2. Identify the interface that you want to act as the upstream interface.
3. Enable IGMP proxy on that interface.

   host1(config-if)#**ip igmp-proxy**

4. (Optional) Specify how often the router sends unsolicited reports to routers on the upstream interface.

host1(config-if)#**ip igmp-proxy unsolicited-report-interval 600**

5. (Optional) Specify how long the router calculates an IGMPv1 querier router to exists on the subnetwork after the router receives an IGMPv1 query on this interface.

   host1(config-if)#**ip igmp-proxy V1-router-present-time 600**

Related Documentation

- IGMP Proxy Overview on page 68
- Establishing the IGMP Proxy Baseline on page 70

Establishing the IGMP Proxy Baseline

You can set the counters for the number of queries received and reports sent on the upstream interface to zero. This feature enables you to establish a reference point, or baseline, for IGMP proxy statistics.

(host1)#**baseline ip igmp-proxy interface**

Related Documentation

- Configuring IGMP Proxy on page 69
- baseline ip igmp-proxy interface

CHAPTER 4

Monitoring IGMP and IGMP Proxy

IP hosts use Internet Group Management Protocol (IGMP) in IPv4 to report their multicast group memberships to neighboring routers. Similarly, multicast routers, such as an E Series router, use IGMP to discover which of their hosts belong to multicast groups.

NOTE: The output for **monitor** and **show** commands from E120 and E320 routers is identical to the output from other E Series routers, except that the output from E120 and E320 routers includes information about the adapter identifier in the interface specifier (slot/adapter/port).

The following topics describe how to monitor IGMP for IP multicast and IGMP proxy on an E Series router:

Monitoring IGMP Information on a Virtual Router

Purpose Display IGMP information for a virtual router.

Action To display IGMP information for a virtual router:

```
host1:boston#show ip igmp
Routing Process IGMP, Administrative state enabled, Operational state enabled
    2 total interfaces, 2 enabled, 0 disabled
0 enabled interfaces performing graceful restart
2 learnt groups
IGMP Statistics:
    Rcvd: 1 total, 0 checksum errors, 0 unknown types
          0 queries, 1 reports, 0 leaves
    Sent: 11 total
```

Meaning Table 16 on page 72 lists the **show ip igmp** command output fields.

Table 16: show ip igmp Output Fields

Field Name	Field Description
Routing Process	Routing process for this virtual router (IGMP)
Administrative state	Status of IGMP in the software: enabled or disabled
Operational state	Status of IGMP on the virtual router: enabled or disabled
Total interfaces	Number of interfaces on which you started IGMP
enabled	Number of interfaces on which IGMP is enabled
disabled	Number of interfaces on which IGMP is disabled
learnt groups	Number of multicast groups that the virtual router has discovered
IGMP graceful restart duration	Restart interval in seconds
IGMP Statistics Rcvd	Statistics for IGMP messages received: • total—Total number of IGMP messages received • checksum errors—Number of IGMP messages received with checksum errors • unknown types—Number of IGMP messages received that are not group membership queries, group membership reports, or leave group membership messages • queries—Number of group membership queries • reports—Number of group membership reports • leaves—Number of leave group membership messages
IGMP Statistics Sent	Statistics for IGMP messages sent: • Total number of group membership queries sent

Related Documentation

- Configuring IGMP Attributes on page 66
- Disabling and Removing IGMP on page 67
- show ip igmp

Monitoring IGMP Groups

Purpose Display statically joined and directly connected groups learned through IGMP.

Action To display statically joined and directly connected groups learned through IGMP without OIF mapping:

```
host1:boston#show ip igmp groups
Grp Address   Interface       State       Reporter     ExpTim oldHTo
------------- --------------- ----------- ------------ ------ ------
228.1.1.      FastEthernet1/1 Version3   17.0.0.2        44      0
228.1.1.2     FastEthernet1/1 Version3   17.0.0.2        50      0
228.1.1.3     FastEthernet1/1 Version3   17.0.0.2        48      0
230.1.1.1     FastEthernet1/1 Version3   17.0.0.2        44      0
  Included Sources:
    51.0.0.1                                                    44
    51.0.0.2                                                    44
    51.0.0.3                                                    44
231.1.1.1     FastEthernet1/1 Version3    17.0.0.2       44      0
  Excluded Sources:
    51.0.0.1                                                     0
    51.0.0.2                                                     0
    51.0.0.3                                                     0

Counts: 5 version-3, 0 version-2, 0 version-1, 0 check state, 0 disabled
       (5 total)
        0 excluded
        Source-groups: 3 included, 3 excluded
```

To display statically joined and directly connected groups learned through IGMP with OIF mapping:

```
host1:boston#show ip igmp groups
Grp Address   Interface       State       Reporter     ExpTim oldHTo
------------- --------------- ----------- ------------ ------ ------
232.1.1.1     ATM5/0.12       Version3    1.1.1.2       371      0
                    oif-map OIFMAP ATM5/0.121
232.1.1.1     ATM5/0.13       Version3    1.1.1.3       375      0
                    oif-map OIFMAP ATM5/0.121
232.1.1.2     ATM5/0.12       Version3    1.1.1.2       373      0
  Included Sources:
    10.1.1.2        oif-map OIFMAP self                        373
    10.1.1.10       oif-map OIFMAP ATM5/0.120                  373
    10.1.1.11       oif-map OIFMAP ATM5/0.121                  373
232.1.1.2     ATM5/0.13       Version3    1.1.1.3       375      0
  Included Sources:
    10.1.1.2        oif-map OIFMAP self                        375
    10.1.1.10       oif-map OIFMAP ATM5/0.120                  375
    10.1.1.11       oif-map OIFMAP ATM5/0.121                  375

Counts: 4 version-3, 0 version-2, 0 version-1, 0 check state, 0 disabled
       (4 total)
        0 excluded
        Source-groups: 6 included, 0 excluded
```


Meaning Table 17 on page 74 lists the output fields of the **show ip igmp groups** command.

Table 17: show ip igmp groups Output Fields

Field Name	Field Description
Grp Address	Address of the multicast group
Interface	Interface that discovered the multicast group
oif-map	Name of the OIF map and the mapped OIF interface, when a group or source has been mapped to an OIF
State	IGMP version on the interface
ExpTim	Time, in seconds, at which the router stops polling for more members of this group
oldHTo	Time at which the router stops polling for more IGMPv1 members of a group. If this value is 0, the interface has received no IGMPv1 reports for the group
Included Sources	Sources included in the multicast group
Excluded Sources	Sources excluded from the multicast group
Counts	Number of source-group mappings by version and state

Related Documentation

- Configuring Multicast Groups for IGMP on page 62
- Specifying Multicast Groups on page 62
- Assigning a Multicast Group to an Interface on page 62
- Configuring Group Outgoing Interface Mapping on page 62

- show ip igmp groups

Monitoring IGMP Interfaces

Purpose Display IGMP information for interfaces on which IGMP is enabled. You can specify the **count** keyword to view the number of IGMP interfaces. You can also specify the **group** address keyword to view the information for interfaces that belongs to that group.

Action To display IGMP information for interfaces on which IGMP is enabled:

```
host1:boston#show ip igmp interface
Interface ATM2/1.15 address 15.0.0.2/255.255.255.0
Administrative state enabled, Operational state enabled
Interface parameters:
    Version 2
    State Querier
```



```
    Query Interval 125 secs, 53 secs before the next query
    Other querier present interval 250 secs
    Maximum response time 100 (in 10ths of a second)
    Last member query interval 10 (in 10ths of a second)
    Robustness 3
    Interface defaults to global promiscuous mode
    No inbound access group
    No inbound access source-group
    No inbound apply-oif-map
    Immediate Leave: disabled
Explicit Host Tracking: enabled
Max-Group limit: No Limit
    Admission-Bandwidth limit: No Limit
    Group Count: 1
Interface statistics:
    Rcvd: 0 reports, 0 leaves, 0 wrong version queries
    Sent: 1 queries
    Groups learnt: 1

Counts: 0 down, 0 init state, 1 querier, 0 non-querier, 1 Total
Counts: 0 down, 0 init state, 1 querier, 0 non-querier, 1 Total
```

To display IGMP information for a specific interface:

```
host1#show ip igmp interface gigabitEthernet 3/0.0
Interface GigabitEthernet3/0.0 address 10.1.1.1/255.255.255.0
Administrative state enabled, Operational state disabled
Interface parameters:
    Version 2
    State Down
    Query Interval 125 secs
    Other querier present interval 250 secs
    Maximum response time 100 (in 10ths of a second)
    Last member query interval 10 (in 10ths of a second)
    Robustness 3
    Interface defaults to global promiscuous mode
    No inbound access group
    No inbound access source-group
    No inbound apply-oif-map
    Immediate Leave: disabled
Explicit Host Tracking: enabled
Max-Group limit: No Limit
    Admission-Bandwidth limit: No Limit
    Group Count: 0
    IOA packet replication gigabitEthernet 3/8.1
Interface statistics:
    Rcvd: 0 reports, 0 leaves, 0 wrong version queries
    Sent: 0 queries
    Groups learnt: 0
```

Meaning Table 18 on page 75 lists the output fields of the **show ip igmp interface** command.

Table 18: show ip igmp interface Output Fields

Field Name	Field Description
Interface	Type of interface and interface specifier. For details about interface types and specifiers, see Interface Types and Specifiers in the *JunosE Command Reference Guide*.

Table 18: show ip igmp interface Output Fields *(continued)*

Field Name	Field Description
address	IP address of the interface
Administrative state	Status of the interface in the software: enabled or disabled
Operational state	Physical status of the interface: enabled or disabled
Version	IGMP version
State	Function of the interface: querier or nonquerier
Query Interval	Time interval in seconds at which this interface sends query messages
Other querier present interval	Time in seconds that the interface waits before declaring itself as the querier
Maximum response time	Time interval, in tenths of a second, during which this interface waits for a host to respond
Last member query interval	Time, in tenths of a second, that this interface waits before sending a new query to a host that sends a group leave message
Robustness	Number of times this interface sends IGMP messages
Interface defaults to global promiscuous mode	Interface uses the setting of the **igmp promiscuous** command to determine whether it accepts IGMP reports from hosts on any subnetwork
Inbound access group	Information about standard IP access lists configured with the **ip igmp access-group** command
No inbound access group	No access list specified
Inbound access source-group	Information about IP access lists configured with the **ip igmp access-source-group** command
No inbound access source-group	No access list specified
Inbound apply-oif-map	Information about OIF maps configured with the **ip igmp apply-oif-map** command
No inbound apply-oif-map	No map name specified
Immediate Leave	Setting of the **ip igmp immediate-leave** command: enabled or disabled

Table 18: show ip igmp interface Output Fields *(continued)*

Field Name	Field Description
Explicit Host Tracking	Setting of the **ip igmp explicit-tracking** command: enabled or disabled
Max-Group limit	Number of IGMP groups that the interface can accept, as configured with the **ip igmp group limit** command
Admission-Bandwidth limit	Value of the admission-bandwidth limit set for an interface that accepts IGMP groups, or No Limit
Group Count	Number of IGMP groups that the interface has accepted
IOA packet replication	Hardware multicast packet replication interface to which egress multicast packets on this interface are redirected
Interface statistics Rcvd	Information about IGMP messages received on this interface: • reports—Number of group membership reports received • leaves—Number of group leave messages received • wrong version queries—Number of group membership queries received from devices running a different version of IGMP
Interface statistics Sent	Number of IGMP messages this interface has sent
Interface statistics Groups learnt	Number of groups this interface has discovered
Counts	Breakdown of IGMP interfaces: • down—Number of interfaces down • init state—Number of interfaces in the initialization state • querier—Number of querier interfaces • non-querier—Number of non-querier interfaces • Total—Total number of IGMP interfaces

Related Documentation

- Configuring IGMP Settings for an Interface on page 61
- Configuring Multicast Groups for IGMP on page 62
- Specifying Multicast Groups on page 62
- Configuring Group Outgoing Interface Mapping on page 62
- Configuring IGMP Attributes on page 66
- Overview of Limiting the Number of Accepted IGMP Groups on page 64

- Configuring Explicit Host Tracking on page 67
- show ip igmp interface

Monitoring Summary Information for IGMP Interfaces

Purpose Display a summary of IGMP information for interfaces on which IGMP is enabled.

Action To display a summary of IGMP information for interfaces on which IGMP is enabled:

```
host1:boston# show ip igmp interface brief
Interface         Intf Address      Ver  State    Querier        QTime  QPTime
--------------    ----------        ---  ------   -------------  -----  ---
fastEthernet0/0   192.168.1.250/24   2   Querier  192.168.1.250  28     0
atm3/0.2          21.1.1.1/8         2   Querier  21.1.1.1       26     0
Count: 2 interfaces
```

Meaning Table 19 on page 78 lists the output fields of the **show ip igmp interface brief** command.

Table 19: show ip igmp interface brief Output Fields

Field Name	Field Description
Interface	Type of interface and interface specifier. For details about interface types and specifiers, see Interface Types and Specifiers in the *JunosE Command Reference Guide*.
Intf Address	IP address of the interface
Ver	IGMP version
State	Function of the interface: querier or nonquerier
Querier	IP address of the querier on the network to which this interface connects
QTime	Time interval, in seconds, at which this interface sends query messages
QPTime	Time in seconds that the interface waits before declaring itself as the querier
Count	Total number of IGMP interfaces

Related Documentation

- Configuring IGMP Settings for an Interface on page 61
- show ip igmp interface

Monitoring IGMP on a Mapped Outgoing Interface

Purpose Display the current mappings to all the mapped outgoing interfaces or to the specified mapped outgoing interface.

Action To display the current mappings to all the mapped outgoing interfaces:

```
host1# show ip igmp mapped-oif
     OIF        Oper  Group Address   Source Address     Join I/F       Map Name
--------------- ---- --------------- --------------- --------------- ----------
ATM5/0.120       Up   232.1.1.2        10.1.1.10       ATM5/0.12       OIFMAP
                                                       ATM5/0.13       OIFMAP
ATM5/0.121       Up   232.1.1.1        *               ATM5/0.12       OIFMAP
                                                       ATM5/0.13       OIFMAP
                      232.1.1.2        10.1.1.11       ATM5/0.12       OIFMAP
                                                       ATM5/0.13       OIFMAP

Counts: 3 source-group mappings
```

Meaning Table 20 on page 79 lists the output fields of the **show ip igmp mapped-oif** command.

Table 20: show ip igmp mapped-oif Output Fields

Field Name	Field Description
OIF	Outgoing interface used in an OIF map
Oper	Operation status of the outgoing interface
Group Address	Multicast group IP address associated with the OIF
Source Address	Source IP address associated with the OIF
Join I/F	IGMP interface associated with the OIF
Map Name	Name of the map associated to the OIF
Counts	Number of source-group mappings to OIFs

Related Documentation

- Configuring Multicast Groups for IGMP on page 62
- Configuring Group Outgoing Interface Mapping on page 62

- show ip igmp mapped-oif

Monitoring IGMP for Multicast Groups

Purpose Display the IGMP membership information for multicast groups and (S, G) channels. You can specify the **tracked** keyword to view the interface information only for interfaces where explicit host tracking is enabled.

Action To display IGMP membership information for multicast groups and (S, G) channels:


```
host1# show ip igmp membership
Flags: M - Uses Oifmap  S- SSM mapped  T - tracked
       1,2,3 - The version of IGMP the group is in
Reporter:
<ip-address> - last reporter if the group is not explicitly tracked
<n>/<m> - <n> reporters include mode, <m> reporters in exclude
Group            Source            Reporter        ExpTim    Flags    Interface
---------------  ---------------   -------------   ------    ------
---------------
224.0.1.40           *             10.10.1.1       02:41     2S
FastEthernet2/1
224.0.1.50                         1/2             02:56     3MT
FastEthernet2/2
                                   11.10.0.21      02:56
                                   11.10.2.22      02:30
                 20.30.0.11
                                   11.10.0.23      02:48
                 20.30.0.12
                                   11.10.0.21      02:56
                 20.30.0.13
                                   11.10.0.21      02:56
                                   11.10.0.22      02:30
                                   11.10.0.23      02:48
224.0.1.60                         20.20.0.1       01:56     3
FastEthernet2/3
                 10.30.0.100                       02:45
                 10.30.0.101                       02:35
                 10.30.0.102                       02:15
                 10.30.0.104                       stop
224.0.1.70                         30.20.0.1       stop      3
FastEthernet2/4
                 40.30.0.100                       01:10
                 40.30.0.101                       01.24
239.0.1.80                         2/0             stop      3T
FastEthernet2/5
                 50.30.0.100
                                   10.10.0.10      02:48
                 50.30.0.101
                                   10.10.0.20      02:56
                                   10.10.0.10      02:48
                 50.30.0.102       10.10.0.20      02:56
235.0.1.90                         0/3             02:56     2T
FastEthernet2/6
                      *
                                   12.10.0.10      02:48
                                   12.10.0.20      02:56
                                   12.10.0.30      02:48
```

Meaning Table 21 on page 80 lists the output fields of the **show ip igmp membership** command.

Table 21: show ip igmp membership Output Fields

Field Name	Field Description
Group	Multicast group or (S, G) channel
Source	(S, G) entries that are forwarding traffic

Table 21: show ip igmp membership Output Fields *(continued)*

Field Name	Field Description
Reporter	Hosts that requested including sources or have not requested excluding sources. If listed under a group, host that sent exclude reports for the group. If listed under a source, host that requested traffic from this source for the group. For any (S, G), if listed under a source, hosts interested in the traffic for this (S, G).
Flags	• M—Uses Oifmap • S—SSM mapped • T—Tracked • 1, 2, 3—IGMP version that the group is in
ExpTim	Expiration time
Interface	Type of interface and interface specifier. For details about interface types and specifiers, see Interface Types and Specifiers in the *JunosE Command Reference Guide*.

Related Documentation

- Configuring Multicast Groups for IGMP on page 62
- Assigning a Multicast Group to an Interface on page 62
- Configuring IGMP Attributes on page 66
- IGMP Traffic Overview on page 65
- Configuring Explicit Host Tracking on page 67
- show ip igmp membership

Monitoring IGMP on Outgoing Interfaces

Purpose Display all the OIF maps or the OIF map for the specified map name.

Action To display all the OIF maps:

```
host1#show ip igmp oif-map
     Map Name           Group Prefix       Source Prefix           OIF
------------------ ------------------ ------------------ ------------------
OIFMAP             232.1.1.0/24       0.0.0.0/0          ATM5/0.121
                   232.1.1.0/24       10.1.1.2/32        self
                   232.1.1.0/24       10.1.1.10/32       ATM5/0.120
                   232.1.1.3/32       0.0.0.0/0          ATM5/0.130
                   232.1.1.4/32       0.0.0.0/0          ATM5/0.130
```

Meaning Table 22 on page 82 lists the output fields of the **show ip igmp oif-map** command.

Table 22: show ip igmp oif-map Output Fields

Field Name	Field Description
Map Name	Name of the map associated to the show output
Group Prefix	Multicast group IP prefix
Source Prefix	Source IP prefix
OIF	Outgoing interface associated with the group and source prefix

Related Documentation

- Configuring Multicast Groups for IGMP on page 62
- Configuring Group Outgoing Interface Mapping on page 62
- show ip igmp oif-map

Monitoring IGMP Information for a Mapped Outgoing Interface

Purpose Display the mapped OIF that is assigned to a given map name, group address, and source address.

Action To display the mapped OIF that is assigned to a given map name, group address, and source address:

`host1#`**show ip igmp oif-mapping OIFMAP 232.1.1.1 10.1.1.10**

```
OIF Mapping
OIF-MAP Name    : OIFMAP
Group Address   : 232.1.1.1
Source Address  : 10.1.1.10
Mapped OIF      : ATM5/0.120
```

Meaning Table 23 on page 82 lists the output fields of the **show ip igmp oif-mapping** command.

Table 23: show ip igmp oif-mapping Output Fields

Field Name	Field Description
OIF-MAP Name	Name of the map requested
Group Address	Multicast group IP address requested
Source Address	Source IP address requested
Mapped OIF	Interface associated with the OIF map

Related Documentation

- Configuring Multicast Groups for IGMP on page 62
- Configuring Group Outgoing Interface Mapping on page 62

- show ip igmp oif-mapping

Monitoring IGMP SSM Mapping

Purpose Display the SSM mapping state and the source list mapping associated with a multicast group address, based on the static SSM mapping configuration.

Action To display the SSM mapping state and the source list mapping associated with a multicast group address:

```
host1:boston#show ip igmp ssm-mapping 232.1.1.1
    SSM Mapping   : Enabled
    Group Address : 232.1.1.1
    Source List   : 172.1.1.1
                  : 172.1.1.2
```

Meaning Table 24 on page 83 lists the output fields of the **show ip igmp ssm-mapping** command.

Table 24: show ip igmp ssm-mapping Output Fields

Field Name	Field Description
SSM Mapping	Status of SSM mapping on the interface: Enabled or Disabled
Group Address	Multicast group address requested
Source List	List of sources mapped to the multicast group address

Related Documentation

- Configuring IGMP Attributes on page 66
- SSM Mapping on page 63
- show ip igmp ssm-mapping

Monitoring the Number of IGMP Groups on a Port

Purpose Display the number of IGMP groups that ports have accepted and, if configured, the maximum number of groups that ports can accept.

NOTE: Only ports that have accepted IGMP groups and ports for which you have configured a limit for the number of IGMP groups appear in this output.

Action To display the number of IGMP groups that ports have accepted and, if configured, the maximum number of groups that ports can accept:

```
host1:boston#show multicast group limit
Port      limit count
```



```
--------- ----- -----
2/0           5     0
2/1          -1     1
```

Meaning Table 25 on page 84 lists the output fields of the **show multicast group limit** command.

Table 25: show multicast group limit Output Fields

Field Name	Field Description
Port	Identifier of the port in *slot/port* format: • slot—Number of the chassis slot in the range 0–6 (ERX7xx models) and 0–13 (ERX14xx models) • port—Port number on the I/O module
limit	Maximum number of IGMP groups that the port can accept. A value of –1 indicates that no limit has been specified.
count	Actual number of IGMP groups that the port has accepted

Related Documentation

- Configuring IGMP Attributes on page 66
- Overview of Limiting the Number of Accepted IGMP Groups on page 64
- show multicast group limit

Monitoring IGMP Proxy Parameters

Purpose Display the IGMP proxy parameters for a virtual router.

Action To display the IGMP proxy parameters for a virtual router:

```
host1#show ip igmp-proxy
Routing Process IGMP Proxy, Administrative state enabled, Operational state enabled
total 1 upstream interface, state enabled
    6 multicast group
```

Meaning Table 26 on page 84 lists the output fields of the **show ip igmp-proxy** command.

Table 26: show ip igmp-proxy Output Fields

Field Name	Field Description
Routing Process	IGMP proxy protocol
Administrative state	State of IGMP proxy in the software: enabled or disabled
Operational state	Operational state of IGMP proxy: enabled or disabled

Table 26: show ip igmp-proxy Output Fields *(continued)*

Field Name	Field Description
total interface	Number of IGMP proxy interfaces on the virtual router; currently only one upstream interface per virtual router
state	Operational state of the IGMP proxy interfaces: enabled or disabled
multicast group	Number of multicast groups associated with IGMP proxy interfaces

Related Documentation

- Configuring IGMP Proxy on page 69
- show ip igmp-proxy

Monitoring IGMP Proxy Groups

Purpose Display information about the multicast groups that IGMP proxy reported.

Action To display information about the multicast groups that IGMP proxy reported:

```
host1#show ip igmp-proxy groups
Grp Address     Interface       Member State
--------------- --------------- --------------
225.1.1.1       atm3/0.2          Idle
225.1.1.2       atm3/0.2          Idle
225.1.1.3       atm3/0.2          Idle
225.1.1.4       atm3/0.2          Idle
225.1.1.5       atm3/0.2          Idle
225.1.1.6       atm3/0.2          Idle
count 6
```

To display information about a particular multicast group:

```
host1#show ip igmp-proxy group 225.1.1.1
Grp Address     Interface       Member State
--------------- --------------- --------------
225.1.1.1       atm3/0.2        Idle
```

To display the number of groups that IGMP proxy reported:

```
host1#show ip igmp-proxy group count
Count: 6 groups
```

Meaning Table 27 on page 85 lists the output fields of the **show ip igmp-proxy groups** command.

Table 27: show ip igmp-proxy groups Output Fields

Field Name	Field Description
Grp Address	Address of the multicast group

Table 27: show ip igmp-proxy groups Output Fields *(continued)*

Field Name	Field Description
Interface	Type and specifier of the upstream interface associated with the multicast group
Member State	State of the associated group address and interface: • Idle—Interface is going to send a group membership report to respond to a group membership query for this group • Delay—Interface has responded to the latest group membership query for this group
count	Total number of multicast groups associated with this interface

Related Documentation

- show ip igmp-proxy groups

Monitoring IGMP Proxy Interfaces

Purpose

Display information about the interface on which IGMP proxy is configured. To view information about a particular interface, enter an interface type and specifier, such as atm 3/0. For details about interface types and specifiers, see Interface Types and Specifiers in the *JunosE Command Reference Guide*. You can specify the **brief** keyword to display a summary rather than a detailed description.

Action

To display information about the specific interface on which IGMP proxy is configured:

```
host1#show ip igmp-proxy interface atm 3/0.2
Interface atm3/0.2 address 21.1.1.1/255.0.0.0
Administrative state enabled, Operational state enabled
Interface parameters:
    Version 2
    State No v1 Router Present
    Unsolicited report interval 10 secs
    Version 1 router present timeout 400 secs
    0 multicast group
Interface statistics:
    Rcvd:  0 v1 query, 6 v2 queries
           0 v1 report, 0 v2 report
    Sent:  0 v1 report, 48 v2 reports, 0 leave
```

Meaning

Table 28 on page 86 lists the output fields of the **show ip igmp-proxy interface** command.

Table 28: show ip igmp-proxy interface Output Fields

Field Name	Field Description
Interface	Type of upstream interface. For details about interface types, see Interface Types and Specifiers in the *JunosE Command Reference Guide*.

Table 28: show ip igmp-proxy interface Output Fields *(continued)*

Field Name	Field Description
address	Address of the upstream interface
Administrative state	State of the upstream interface in the software: enabled or disabled
Operational state	Physical state of the upstream interface: enabled or disabled
Version	IGMP version on this interface
State	Presence of IGMPv1 routers on the same subnet as this upstream interface
Unsolicited report interval	Time interval, in tenths of a second, at which this upstream interface sends an unsolicited group membership report
Version 1 router present timeout	How long, in seconds, that the upstream interface calculates an IGMPv1 router to exist on the subnet after that interface receives an IGMPv1 group membership query
multicast group	Number of multicast groups associated with this upstream interface
Interface statistics Rcvd	Statistics for messages received on this interface: • v1 queries—Number of IGMPv1 group membership queries received • v2 queries—Number of IGMPv2 group membership queries received • v1 reports—Number of IGMPv1 group membership reports received • v2 reports—Number of IGMPv2 group membership reports received
Interface statistics Sent	Statistics for messages sent from this interface: • v1 reports—Number of IGMPv1 leave group reports sent • v2 reports—Number of IGMPv2 leave group reports sent • leaves—Number of leave group membership messages sent

Related Documentation

- Configuring IGMP Proxy on page 69
- Establishing the IGMP Proxy Baseline on page 70
- show ip igmp-proxy interface

CHAPTER 5

Configuring PIM for IPv4 Multicast

The Protocol Independent Multicast (PIM) protocol is a collection of multicast routing protocols that enables multicast routers to identify other multicast routers to receive packets.

This chapter describes how to configure PIM for IPv4 on E Series routers; it contains the following sections:

Understanding PIM for IPv4 Multicast

The IPv4 implementation of PIM supports PIM dense mode, PIM sparse mode, PIM sparse-dense mode, and PIM source-specific multicast (PIM SSM).

Figure 8 on page 90 represents how PIM builds a source, group (S,G) entry in a source-rooted tree (SRT). When multiple routers are connected to a multiaccess network, one router becomes the designated router. The designated router receives data from the source on interface 1/0 and multicasts the data to its downstream neighbors on interfaces 1/1, 2/0, and 2/1. In the designated router routing table, the entry for this operation lists the source as the IP address of the source and the group as the IP address of the multicast group.

Figure 8: Source-Rooted Tree

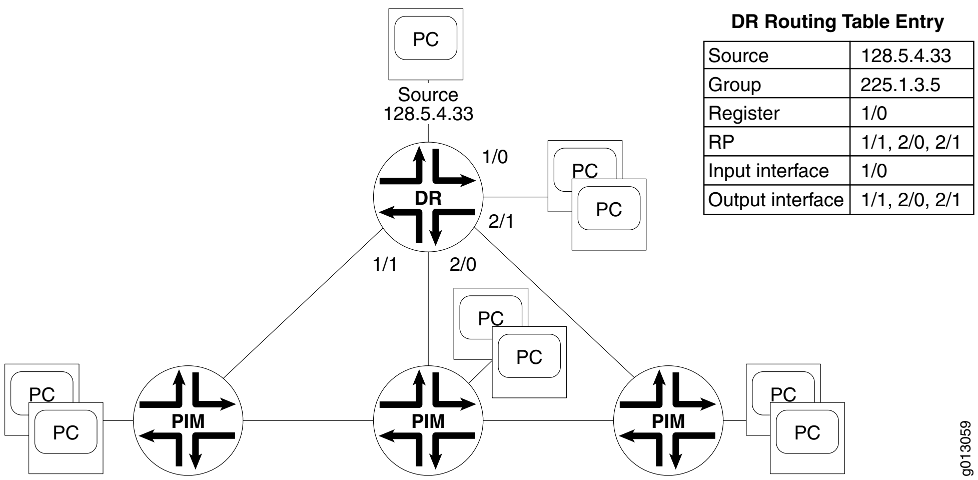

Neighbors exchange hello messages periodically to determine the designated router. The router with the highest network layer address becomes the designated router. If the designated router subsequently receives a hello message from a neighbor with a higher network layer address, that neighbor becomes the designated router.

The IPv4 implementation of PIM supports the following modes:

PIM Dense Mode

PIM dense mode uses a reverse-path multicast, flood-and-prune mechanism. The protocol was developed for situations that meet one or more of the following criteria:

- Sources and receivers are close together, and there are many more receivers than sources.

- There is a constant stream of multicast data.
- There is a lot of multicast data.

Dense-mode routing protocols use *SRT algorithms*. An SRT algorithm establishes a tree that connects each source in a multicast group to the members of the group. All traffic for the multicast group passes along this tree.

Figure 9 on page 91 illustrates how PIM dense mode works. When a source sends a multicast packet to a first-hop router, the first-hop router multicasts that packet to its neighbors. Those neighbors in turn forward the packet to their neighbors and their hosts that belong to the multicast group. If a neighbor has no hosts that belong to the multicast group and has no other PIM neighbors, it returns a prune message to the first-hop router. The first-hop router does not multicast subsequent packets for that group to neighbors who respond with prune messages.

Figure 9: PIM Dense Mode Operation

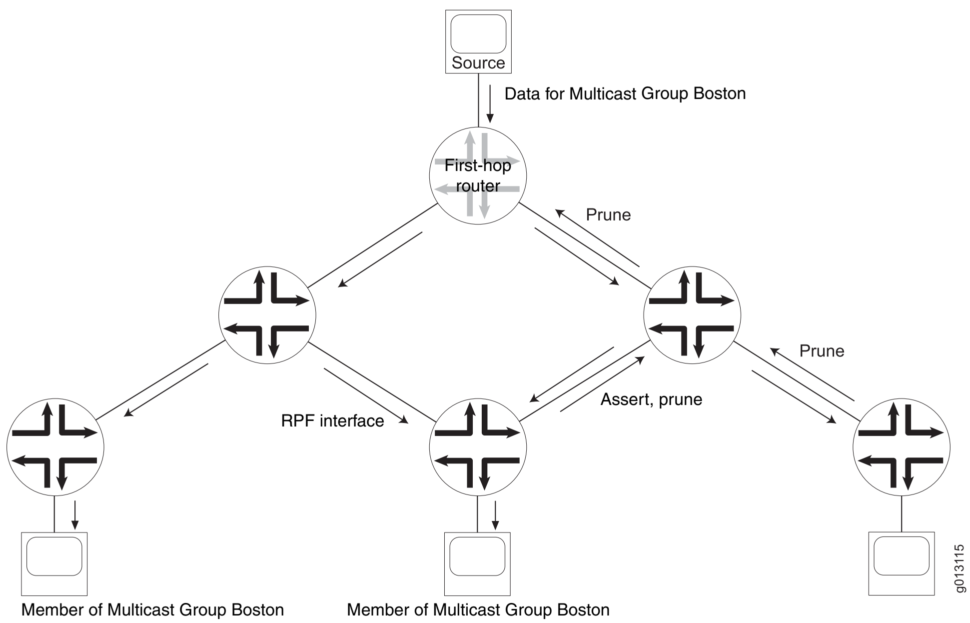

Overriding Prunes

If a host on a previously pruned branch requests to join a multicast group, it sends an IGMP message to its first-hop router. The first-hop router then sends a graft message upstream.

PIM routers send join messages on multiaccess interfaces to override prune messages. For example, if a PIM router sent a prune message to indicate that it had no hosts for a multicast group, and one of its hosts subsequently requests to send a packet to that group, the router sends a join message to the first-hop router.

Preventing Duplication

If there are parallel paths to a source, duplicate packets can travel downstream through different routers to the network. If a forwarding router receives a multicast packet on its outgoing interface, the router identifies that the packet is a duplicate and notifies the upstream routers. See Figure 10 on page 92.

Figure 10: Detecting Duplication

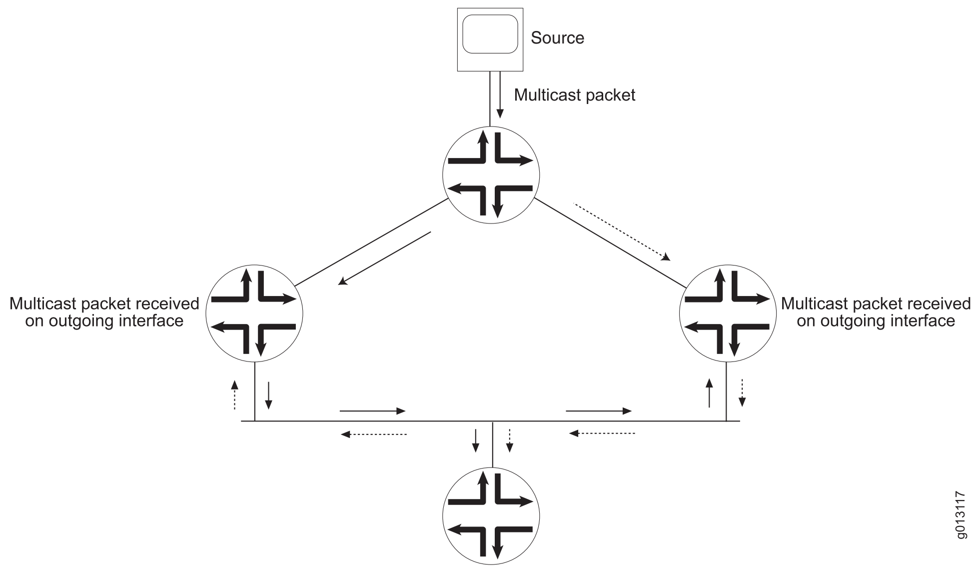

The upstream routers responsible for the duplication send assert messages to determine which router becomes the forwarder. Downstream routers listen to the assert messages to discover which router becomes the forwarder.

PIM Sparse Mode

This implementation of PIM sparse mode supports the following features:

- Rendezvous point (RP) routers
- Designated routers and designated router election
- Join/prune messages, hello messages, assert messages, and register messages
- Switching from a shared tree to a shortest path tree (SPT)
- (*,*,RP) support for interoperation with dense-mode protocols
- RPF checks of multicast entries when unicast routing configuration changes
- Timers for tree maintenance
- Border, null, Rendezvous Point Tree (RPT), SPT, and wildcard flags

PIM sparse mode resolves situations that meet one or more of the following criteria:

- The multicast group contains few receivers.
- Multicast traffic is infrequent.
- Wide area networks (WANs) separate sources and receivers.

Sparse-mode routing protocols use *shared trees*. In a shared tree, sources forward multicast datagrams to a directly connected router, the designated router. The designated router encapsulates the datagram and unicasts it to an assigned RP router, which then forwards the datagram to members of multicast groups. See Figure 11 on page 93.

Figure 11: PIM Sparse Mode Operation

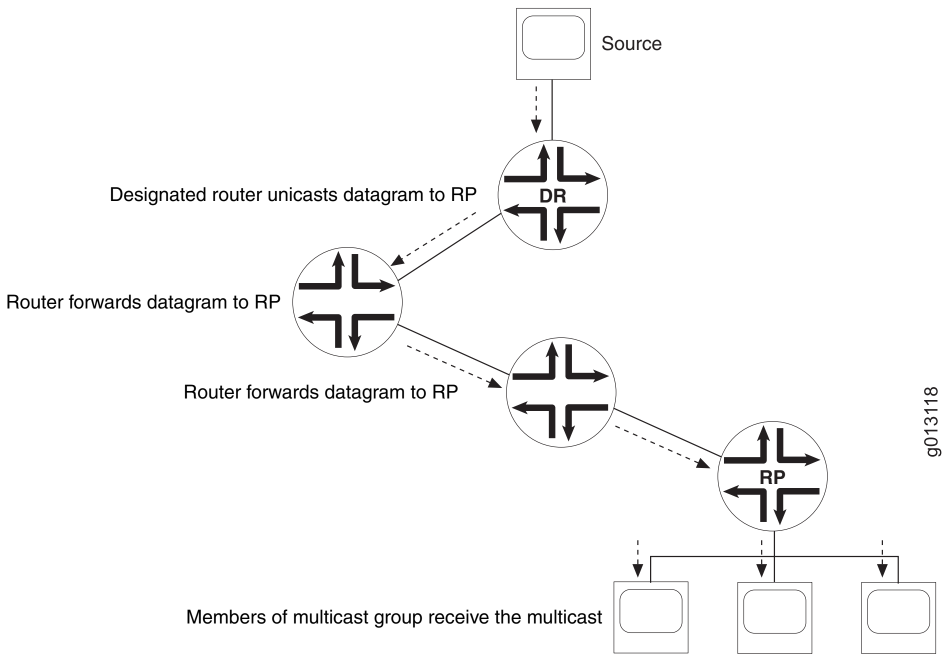

In PIM sparse mode, an RP announces a source and establishes paths from the source to members of a multicast group before multicasting any datagrams. RPs transmit join messages to become part of the shared tree that enables distribution of packets to the multicast group.

However, when a source starts multicasting datagrams, PIM sparse mode can switch to an SRT—known in PIM sparse mode as an SPT—to improve the network's efficiency. Although shared trees minimize the traffic in the network and the costs associated with unnecessary transmission of data, the routes in a shared tree might be longer than those in an SPT. See Figure 12 on page 94.

Figure 12: Shared Tree Versus SPT

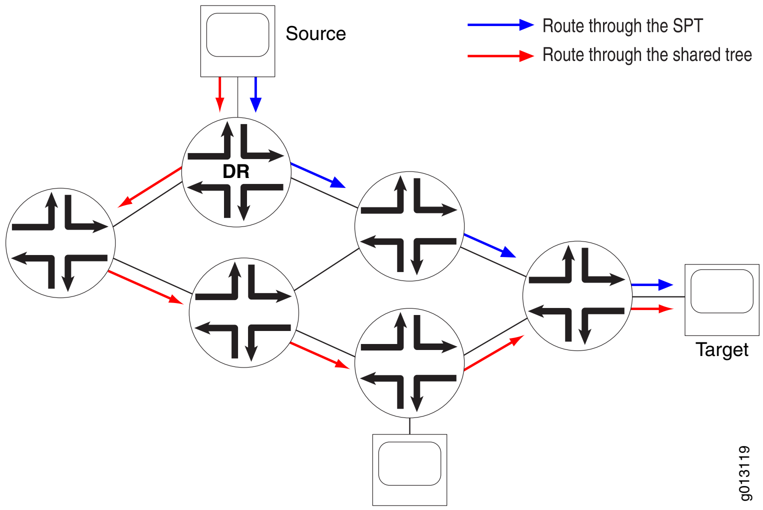

The designated routers on the network determine when the source switches from a shared tree to an SPT. A designated router switches to the SPT when it receives a certain number of packets which you can configure.

When all designated routers associated with a specific RP router have switched to the SPT, the RP router sends a join/prune message toward the multicast source. When the multicast source receives this message, it stops sending multicast data through the SPT.

Joining Groups

A host's designated router (DR) sends join messages to the RP when that host wants to join a group. When a host wants to leave a group, it communicates with its designated router through IGMP. When the designated router no longer has any hosts that belong to a particular group, it sends a prune message to the RP.

Timers

PIM sparse mode uses timers to maintain the networking trees.

NOTE: **PIM sparse mode routers poll their neighbors and hosts for various pieces of information at set intervals.**

If a PIM sparse mode router does not receive information from a neighbor or host within a specific time, known as the *hold time*, it removes the associated information from its routing tables.

You can configure how often an interface sends hello messages (hello interval) and how often routers send RP announce messages (RP announce interval). The hold-time associated with hello messages is 3.5 times the hello interval, and the holdtime associated with RP announce messages is 2.5 times the RP announce interval.

All other timers are fixed and take the default values recommended in RFC 2934—Protocol Independent Multicast MIB for IPv4 (October 2000).

PIM Sparse Mode Bootstrap Router

PIM sparse mode routers need the address of the rendezvous point (RP) for each group for which they have (*,G) state. They obtain this address either through a bootstrap mechanism or through static configuration. PIM sparse mode routers commonly use one of two bootstrap mechanisms: bootstrap router (BSR) or auto-RP. Auto-RP is standards based, but is not used in IPv6 implementations, so BSR configuration has become more popular.

When implemented, BSR operates as follows:

1. One router in each PIM domain is elected the BSR.
2. All the routers in the domain that are configured to be RP candidates periodically unicast their candidacy to the BSR.
3. The BSR picks an RP set from the available candidates and periodically announces this set in a bootstrap message.
4. Bootstrap messages are flooded hop by hop throughout the domain until all routers in the domain learn the RP set.

NOTE: A PIM router can receive group-to-RP mappings from either BSR or auto-RP, but not from both. Because BSR and auto-RP use different mapping algorithms, the mechanisms cannot coexist.

NOTE: Static-override is configured using the override switch in the **ip pim rp-address** command.

CAUTION: The maps distributed by BSR and Auto-RP for the same RP announcements may be different. Because the Auto-RP mapping agent resolves mapping conflicts, a PIM router that applies the BSR mapping algorithm to an Auto-RP distributed map should produce the same result as a router that applies the Auto-RP mapping algorithm. However, the reverse is not true. A PIM router that applies the BSR mapping algorithm to a BSR distributed map may produce a different result to a router that applies the Auto-RP mapping algorithm. This means that **a PIM IPv4 domain can operate either BSR or Auto-RP**.

PIM Sparse-Dense Mode

In PIM sparse-dense mode, if an RP is not known for a group, the router sends data using PIM dense mode. However, if the router discovers an RP or you configure an RP statically, PIM sparse mode takes over.

You can configure both PIM dense mode and PIM sparse mode commands in PIM sparse-dense mode.

PIM Source-Specific Multicast

PIM SSM is an extension of the PIM protocol. Using SSM, a client can receive multicast traffic directly from the source. PIM SSM uses PIM sparse mode functionality to create an SPT between the client and the source, but builds the SPT without using an RP.

By default, the SSM group multicast address is limited to the IP address range 232.0.0.0 to 232.255.255.255. You can use the **ip pim ssm** command to extend SSM operations into another Class D range. (See "Configuring PIM for IPv4 SSM" on page 116.)

An SSM-configured network has the following advantages over a traditionally configured PIM sparse mode network include the following:

- No need for shared trees or RP mapping (no RP is required).
- No need for RP-to-RP source discovery through Multicast Source Discovery Protocol (MSDP).
- Simplified administrative deployment; you need only configure PIM sparse mode on all router interfaces and issue the necessary SSM commands (including specifying IGMPv3 on the receiver local area network).
- Support for source lists; you can use source lists, supported in IGMPv3, where only specified sources send traffic to the SSM group.

In a PIM SSM–configured network, an E Series router subscribes to an SSM channel (by means of IGMPv3 or by means of IGMP ssm-mapping for IGMPv2/v1 joins), requesting to join group G and source S. The directly connected PIM sparse mode router, the designated router of the receiver, sends an (S,G) join message to its RPF neighbor for the source. For PIM SSM, the RP is not contacted in this process by the receiver (as happens in normal PIM sparse mode operations).

Related Documentation

- PIM for IPv4 Multicast Platform Considerations on page 96
- PIM for IPv4 Multicast References on page 97
- Enabling PIM for IPv4 on a Virtual Router on page 97
- Disabling PIM for IPv4 on a Virtual Router on page 98
- Enabling PIM for IPv4 on an Interface on page 98
- Monitoring PIM for IPv4 Multicast on page 121

PIM for IPv4 Multicast Platform Considerations

For information about modules that support PIM on the ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.

- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support PIM.

For information about modules that support PIM on the E120 and E320 Broadband Services Routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support PIM.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- PIM for IPv4 Multicast References on page 97
- Monitoring PIM for IPv4 Multicast on page 121

PIM for IPv4 Multicast References

For more information about PIM, see the following resources:

- Protocol Independent Multicast MIB for IPv4—draft-ietf-idmr-pim-mib-10.txt (July 2000 expiration)
- RFC 2362—Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification (June 1998)
- RFC 3569—An Overview of Source-Specific Multicast (SSM) (July 2003)
- Source-Specific Multicast for IP—draft-ietf-ssm-arch-06.txt (March 2005 expiration)
- Source-Specific Protocol Independent Multicast in 232/8—draft-ietf-mboned-ssm232-08.txt (September 2004 expiration)
- Multicast in MPLS/BGP VPNs—draft-rosen-vpn-mcast-06.txt (April 2004 expiration)
- Multicast in MPLS/BGP IP VPNs—draft-rosen-vpn-mcast-08.txt (June 2005 expiration)

NOTE: **IETF drafts are valid for only 6 months from the date of issuance. They must be considered as works in progress. Please refer to the IETF Web site at http://www.ietf.org for the latest drafts.**

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- PIM for IPv4 Multicast Platform Considerations on page 96
- Monitoring PIM for IPv4 Multicast on page 121

Enabling PIM for IPv4 on a Virtual Router

By default, PIM is disabled. To enable PIM on a VR:

1. Enable multicast routing. (See "Enabling IP Multicast" on page 7.)
2. Create a VR, or access an existing VR context.

    host1(config)#**virtual-router boston**

3. Create and enable PIM processing.

    host1:boston(config)#**router pim**

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- PIM for IPv4 Multicast Platform Considerations on page 96
- PIM for IPv4 Multicast References on page 97
- Disabling PIM for IPv4 on a Virtual Router on page 98
- Removing PIM for IPv4 on page 118
- Resetting PIM Counters and Mappings for IPv4 on page 119
- **virtual-router**
- **router pim**

Disabling PIM for IPv4 on a Virtual Router

To disable PIM processing on a router, use the **pim disable** command. By default, PIM processing is enabled:

- Issue the **pim disable** command in Router Configuration mode.

    host1:boston(config-router)#**pim disable**

  Use the **no** version to reenable PIM processing.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- PIM for IPv4 Multicast Platform Considerations on page 96
- PIM for IPv4 Multicast References on page 97
- Enabling PIM for IPv4 on a Virtual Router on page 97
- Removing PIM for IPv4 on page 118
- Resetting PIM Counters and Mappings for IPv4 on page 119
- **pim disable**

Enabling PIM for IPv4 on an Interface

You can enable PIM on an interface in one of the PIM modes (dense, sparse, or sparse-dense) and specify how often the interface sends hello messages to neighbors.

You can configure PIM and IGMP on the same interface. If you configure IGMP and PIM on an interface, the router determines that PIM owns the interface.

NOTE: You cannot configure DVMRP and PIM on the same interface.

To enable PIM on an interface:

1. Enable PIM on an interface in one of the PIM modes. By default, PIM is enabled in dense mode.

   host1(config-if)#**ip pim sparse-dense-mode**

2. Specify the interval, in seconds, at which the router sends hello messages to neighbors.

   host1(config-if)#**ip pim query-interval 100**

3. Set the graceful restart duration for IP PIM sparse mode.

   host1(config-if)#**ip pim sparse-mode graceful-restart-duration 10**

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- PIM for IPv4 Multicast Platform Considerations on page 96
- PIM for IPv4 Multicast References on page 97
- Enabling PIM for IPv4 on a Virtual Router on page 97
- Removing PIM for IPv4 on page 118
- Resetting PIM Counters and Mappings for IPv4 on page 119

- **ip pim**
- **ip pim query-interval**
- **ip pim sparse-mode graceful-restart-duration**

Setting a Priority to Determine the Designated Router for IPv4

You can influence whether a particular router is selected as the designated router with the **ip pim dr-priority** command. A higher priority value increases the likelihood that a router is selected as the designated router, while a lower value decreases the likelihood. The **ip pim dr-priority** command in Router Configuration mode sets the designated router priority on all the PIM interfaces on the router. To override this global setting on a particular interface, use the **ip pim dr-priority** command in Interface Configuration mode.

NOTE: You cannot set the designated router priority on PIM dense mode interfaces.

- To set a priority value, in the range 1–254, by which a router is likely to be selected as the designated router, issue the **ip pim dr-priority** command.

  In Router Configuration mode:

  host1(config-router)#**ip pim dr-priority 24**

The **no** version restores the default value. The default value is 1.

In Interface Configuration mode:

host1(config-subif)#ip pim dr-priority 24

The **no** version restores the value that is specified in Router Configuration mode.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- PIM for IPv4 Multicast Platform Considerations on page 96
- PIM for IPv4 Multicast References on page 97
- Enabling PIM for IPv4 on a Virtual Router on page 97
- Disabling PIM for IPv4 on a Virtual Router on page 98
- Monitoring PIM for IPv4 Multicast on page 121

- **ip pim dr-priority**

Configuring the PIM for IPv4 Join/Prune Message Interval

When you use the router for PIM, the router sends join/prune messages to the upstream RPF neighbor. The default join/prune message interval is 60 seconds. You can configure the join/prune message interval using the **ip pim join-prune-interval** command. The **ip pim join-prune-interval** command in Router Configuration mode configures the join-prune interval on all the PIM interfaces on the router. To override this global setting on a particular interface, use the **ip pim join-prune-interval** command in Interface Configuration mode.

The hold-time associated with the PIM join/prune interval messages is 3.5 times the PIM join/prune message interval.

NOTE: **You cannot configure the PIM join/prune message interval on PIM DM interfaces.**

- To set an interval value, in the range 10–210 seconds, at which the router sends the PIM join/prune message, issue the **ip pim join-prune-interval** command.

  In Router Configuration mode:

  host1(config-router)#ip pim join-prune-interval 150

  The **no** version restores the default value, 60 seconds.

  In Interface Configuration mode:

  host1(config-subif)#ip pim join-prune-interval 150

  The **no** version restores the value that is specified in Router Configuration mode.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Enabling PIM for IPv4 on a Virtual Router on page 97
- Enabling PIM for IPv4 on an Interface on page 98
- Monitoring PIM for IPv4 Multicast on page 121
- **ip pim join-prune-interval**

Configuring an RP Router for PIM Sparse Mode and PIM Sparse-Dense Mode for IPv4

When you use the router for PIM sparse mode or PIM sparse-dense mode, some VRs must act as RP routers. You can configure static RP routers or configure the router to assign RP routers automatically.

To configure the router to assign RP routers automatically, you must define several VRs as RP routers and one VR as an RP mapping agent. RP routers send their announcement messages to the RP mapping agent, which assigns groups to RP routers and resolves any conflicts. The RP mapping agent notifies neighbors of the RP assigned to each group.

NOTE: You can configure PIM on IPv4 and IPv6 interfaces. However, IPv6 does not support all PIM configuration options. For information about configuring PIM on IPv6 interfaces, see "Configuring PIM for IPv6 Multicast" on page 273.

The following subsections explain how to configure an RP router for PIM sparse mode and PIM sparse-dense mode:

- Configuring a Static RP Router on page 101
- Configuring an Auto-RP Router for PIM Sparse Mode on page 101
- Configuring an Auto-RP Router for PIM Sparse-Dense Mode on page 102

Configuring a Static RP Router

If you want to control PIM more tightly, you can configure a static RP router. To do so:

1. Configure an access list that specifies the multicast groups that can use the static RP router.

   host1(config)#**access-list boston permit 228.0.0.0 15.255.255.255**

2. Specify a static RP router.

   host1(config)#**ip pim rp-address 122.0.0.1 1 boston**

Configuring an Auto-RP Router for PIM Sparse Mode

Two multicast groups, 224.0.1.39 and 224.0.1.40, are reserved for forwarding auto-RP messages through the network. When you configure an auto-RP router for PIM sparse mode, you must assign a static RP router to these two groups. You can then specify an RP mapping agent for other multicast groups.

To configure an auto-RP router for PIM sparse mode:

1. Configure a static RP to have priority over the auto-RP for the groups that send auto-RP multicast messages.

   host1(config)#**access-list 11 permit 224.0.1.39 0.0.0.0**
   host1(config)#**access-list 11 permit 224.0.1.40 0.0.0.0**
   host1(config)#**ip pim rp-address 192.48.1.22 11 override**

2. Assign an RP mapping agent.

   host1(config)#**ip pim send-rp-discovery scope 23 loopback 1**

3. Configure routers to send auto-RP announcement messages to the mapping agent.

   host1(config)#**ip pim send-rp-announce loopback 2 scope 16 group-list 1**

Configuring an Auto-RP Router for PIM Sparse-Dense Mode

In PIM sparse-dense mode, you must prevent routers from advertising auto-RP messages to the multicast groups 224.0.1.39 and 224.0.1.40, which are reserved for forwarding auto-RP messages through the network. To configure an auto-RP router for PIM sparse-dense mode:

1. Assign an RP mapping agent.

   host1(config)#**ip pim send-rp-discovery scope 23 loopback 1**

2. Configure an access list that details the multicast groups that can use the static RP router.

   host1(config)**access-list boston permit 224.0.0.0 15.255.255.255**

3. Prevent routers from advertising auto-RP messages to the multicast groups that are reserved for forwarding auto-RP messages through the network.

   host1(config)#**access-list 1 deny 224.0.1.39**
   host1(config)#**access-list 1 deny 224.0.1.40**

4. Configure routers to send auto-RP announcement messages to the mapping agent.

   host1(config)#**ip pim send-rp-announce loopback 2 scope 23 group-list boston interval 200**

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Configuring the PIM for IPv4 Join/Prune Message Interval on page 100
- Configuring BSR and RP Candidates for PIM Sparse Mode for IPv4 on page 103
- Monitoring PIM for IPv4 Multicast on page 121
- **access-list**
- **ip pim rp-address**
- **ip pim send-rp-discovery scope**

- **ip pim send-rp-announce**

Configuring BSR and RP Candidates for PIM Sparse Mode for IPv4

When choosing candidate BSRs, select well-connected routers in the core of the network. Typically, candidate BSRs are a subset of the candidate RPs. A single BSR is elected for the domain of candidate BSRs. The elected BSR floods bootstrap messages (BSMs) containing their group-to-RP mappings to all PIM routers. PIM routers use the group-to-RP mappings supplied by the elected (or preferred) BSR.

- Issue the **ip pim bsr-candidate** command in Global Configuration mode to define a router as a BSR candidate:

  ```
  host1(config)#ip pim bsr-candidate loopback 1 30 10
  ```

  The **no** version disables the router BSR candidacy.

- Issue the **ip pim rp-candidate** command in Global Configuration mode to define a router as an RP candidate:

  ```
  host1(config)#access-list 1 permit 227.0.0.0 15.255.255.255
  host1(config)#access-list 1 permit 228.0.0.0 15.255.255.255
  host1(config)#ip pim rp-candidate loopback 1 group-list 1
  ```

  The **no** version stops the router from being an RP candidate.

NOTE: You can configure PIM on IPv4 and IPv6 interfaces. However, IPv6 does not support all PIM configuration options. For information about configuring PIM on IPv6 interfaces, see "Configuring PIM for IPv6 Multicast" on page 273.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Monitoring PIM for IPv4 Multicast on page 121
- ip pim rp-address
- ip pim send-rp-announce
- ip pim send-rp-discovery scope

Migrating to BSR from Auto-RP

Migrating to BSR from auto-RP requires that you upgrade all PIM routers in the domain to support BSR. However, until all routers are BSR-capable, continue to use auto-RP.

After all routers are BSR-capable, switch from auto-RP to BSR as follows:

1. Use the **no ip pim send-rp-discovery scope** command to stop PIM in the network by disabling all auto-RP mapping agents. This results in flooding to an empty map.

2. Reconfigure auto-RP mapping agents as candidate BSRs by using the **ip pim bsr-candidate** command.
3. Reconfigure auto-RP candidate RPs as BSR candidate RPs by issuing the **no ip pim send-rp-announce** command and then issuing the **ip pim rp-candidate** command.

NOTE: You can configure PIM on IPv4 and IPv6 interfaces. However, IPv6 does not support all PIM configuration options. For information about configuring PIM on IPv6 interfaces, see "Configuring PIM for IPv6 Multicast" on page 273.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Monitoring PIM for IPv4 Multicast on page 121
- **ip pim rp-candidate**
- **ip pim bsr-candidate**

Switching to an SPT for PIM Sparse Mode for IPv4

PIM sparse mode initiates multicasting using a shared tree. You can configure PIM sparse mode to switch to an SPT when a source starts sending multicast messages, or you can prevent PIM sparse mode from switching to an SPT. Multicasting over an SPT might be more efficient than multicasting over a shared tree. (See *PIM Sparse Mode* in "Understanding PIM for IPv4 Multicast" on page 90.)

To specify the network configuration that PIM sparse mode uses when a source starts sending multicast messages:

- Issue the **ip pim spt-threshold** command in Global Configuration mode.

  ```
  host1(config)#ip pim spt-threshold
  ```

  The **no** version restores the default value, 0.

NOTE: You can configure PIM on IPv4 and IPv6 interfaces. However, IPv6 does not support all PIM configuration options. For information about configuring PIM on IPv6 interfaces, see "Configuring PIM for IPv6 Multicast" on page 273.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Monitoring PIM for IPv4 Multicast on page 121
- **ip pim spt-threshold**

Multicast VPNs Overview

JunosE Software provides the ability to create multicast VPNs by using GRE tunnels. This implementation is based on *Multicast in MPLS/BGP VPNs* (draft-rosen-vpn-mcast-06.txt and draft-rosen-vpn-mcast-08.txt) and further defined by *Base Specification for Multicast in MPLS/BGP VPNs* (draft-raggarwa-13vpn-2547-mvpn-00.txt).

NOTE: **Although you can configure PIM sparse mode remote neighbors, you can no longer use these remote neighbors for BGP/MPLS VPNs. For multicast VPNs, use the functionality described in this section.**

The JunosE Software supports default Multicast Distribution Trees (MDTs) and data MDTs. The following topics explain how to create multicast VPNs using the default MDTs and the data MDTs:

- Creating Multicast VPNs Using the Default MDT on page 105
- Creating Multicast VPNs Using the Data MDT on page 105

Creating Multicast VPNs Using the Default MDT

The JunosE Software does not support a single MDT command. Instead, you must configure the multicast tunnel interfaces (MTIs) explicitly. The MTI is an IP interface that is stacked on a GRE tunnel interface. The destination address of the GRE tunnel is the multicast VPN (MVPN) group address of the MDT.

A **tunnel mdt** command specifies that the tunnel is the MTI for the default MDT, enabling the creation of a second, layer 2 interface (interface tunnel gre:*name*.mdt) on which an unnumbered IP interface (tied to the provider edge loopback interface) is stacked in the context of the parent virtual router.

Creating Multicast VPNs Using the Data MDT

A data multicast distribution tree (MDT), based on section 8 of Internet draft draft-rosen-vpn-mcast-08.txt, *Multicast in MPLS/BGP IP VPNs*, solves the problem of P routers flooding unnecessary multicast information to PE routers that have no interested receivers for a particular VPN multicast group. The data MDT solution requires the creation of a new tunnel by the PE router if the source exceeds a configured rate threshold parameter. All other PE routers join the new tunnel only if the PE router has receivers in the VPN for that multicast group.

The JunosE Software uses dynamic point-to-multipoint GRE tunnels to configure data MDTs. In the current release, IPv6 transport over GRE (unicast or multicast) is not supported. For more information, see *Configuring Dynamic IP Tunnels* in the *JunosE IP Services Configuration Guide*.

Data MDTs are established using PIM-SM (shared RP Trees) and PIM-SSM (Source Trees). Profiles for dynamic interfaces in the VRF are restricted to sparse-mode only.

Data MDT Sources

A C-SG flow arriving in the source VRF is a candidate for a data MDT if the system matches the C-SG in the route map that you specify for the data MDT using the **ip pim data-mdt** command. The C-SG flow is initially forwarded on the default MDT. The system creates the data MDT when the flow rate exceeds a value you configure in the route map using the **set threshold** command.

When the Source C-PIM-SM first creates a data MDT for a C-SG flow, it sends a <C-SG, P-G> MDT join message with type, length, value (TLV) format to the default MDT. This message invites peer PE routers to join the new data MDT. It starts a timer that you can configure using the **mdt-data-delay** command to track the number of seconds before switching to the data MDT. When that timer expires, C-PIM-SM switches from sending C-SG data on the default MDT to sending data on the data MDT.

When the C-SG flow is switched to the data MDT, the Source C-PIM-SM starts a timer that you can configure using the **mdt-data-holddown** command to track the number of seconds before switching to the default MDT. When the timer expires, the data MDT is deleted and the C-SG flow switched back to the default MDT if the flow rate drops back below the threshold. If the flow rate exceeds the threshold, the timer restarts. If the timer expires and the flow rate is below the threshold, the data MDT is removed.

The Source C-PIM-SM maintains sent MDT Join TLV messages in its database as long as they are active. While the data MDT is active, C-PIM-SM resends that MLD Join TLV message using a setting that you can configure using the **mdt-interval** command to measure time in seconds between successive MLD join TLV messages.

Data MDT Receivers

When the Receiver C-PIM-SM receives a <C-SG, P-G> MDT Join TLV message from the default MDT, it extracts the C-SG and the data MDT P-Group address from the TLV and queries the route map that you specified for the data MDT to determine whether the C-SG is a candidate for a data MDT. If it matches, the C-PIM-SM adds the MDT Join TLV to its database and records the time.

If the Receiver C-PIM-SM does not receive an MDT Join TLV<C-SG, P-G> to refresh its database within the amount of time specified for the timeout in the **mdt-data-timeout** command, the MDT Join TLV<C-SG> is removed from the database and the associated data MDT is removed.

When a new MDT Join TLV<C-SG, P-G> is added to the database, the Receiver C-PIM-SM determines whether it has an SG, SPT state. If it has an SG state, and the incoming interface (IIF) is the default MDT, then C-PIM-SM creates the data MDT and deletes the corresponding forwarding entry. C-PIM-SM waits for the source to transmit data on the data MDT. During this period, data can continue to be received on the default MDT. C-PIM-SM fails the reverse-path forwarding (RPF) check, which results in a forwarding entry with a discarded IIF.

If the C-SG,SPT state is created (either as a result of a C-SSM join or switch from RPT to SPT), and it is the default MDT, the Receiver C-PIM-SM determines whether an MDT Join TLV<C-SG> is active. If it is, C-PIM-SM creates the data MDT.

Establishing a Data MDT Using ASM or SSM

A data MDT carries one C-SG flow. If the data MDTs are established using any-source multicast (ASM), then the P-Group address selected by a PE for the data MDT must be unique to that PE in the MDT (that is, the range of MDT P-Group addresses available in the core must be administratively divided among all the PEs that will source VPN multicasts). The VRFs in a PE must share the P-Group addresses in the assigned range for the PE.

If the data MDTs are established using single-source multicast (SSM), you must configure VRFs to transmit on a tunnel using the same MDT P-Group address. Each VRF transmits using a unique P-Source address; however, each data MDT created by the VRF must use a different P-Group address. There might be one sender data MDT and possibly many receiver data MDTs sharing an IP tunnel. Each PE can assign MDT P-Groups from the same range, but the P-Group addresses must be administratively divided among the VPNs.

For a receiver on the data MDT, P-PIM-SM joins the data MDT by propagating join state into the core. The P-Group for that join is extracted from the MDT Join TLV. If SSM is not activated or the P-Group is not in the SSM group range, P-PIM-SM performs a <*, G> join towards the RP for that P-Group.

If SSM is activated and the P-Group is in the SSM group range, P-PIM-SM performs an <S, G> join towards the P-Source, where the P-Source address is the SA of the MDT Join TLV.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Configuring the Default MDT on page 107
- Configuring Data MDTs on page 109
- Example: Configuring Multicast VPNs on page 111
- Monitoring PIM for IPv4 Multicast on page 121

Configuring the Default MDT

To configure the default MDT:

1. On the parent virtual-router, configure an IP interface over the MDT interface, and make it a PIM interface.

```
host1(config)#virtual-router pe1
host1:pe1(config)#ip multicast-routing
host1:pe1(config)#interface loopback 0
host1:pe1(config-if)#ip address 1.1.1.1 255.255.255.255
host1:pe1(config-if)#ip pim sparse-mode
host1:pe1(config-if)#exit
```

2. Configure the VRF.

```
host1(config)#virtual-router pe1
host1:pe1(config)#ip vrf pe13
```


host1:pe1(config-vrf)#**rd 100:13**
host1:pe1(config-vrf)#**route-target both 100:3**
host1:pe1(config-vrf)#**exit**

3. To enable MDT, configure a GRE multicast tunnel interface (MTI) for VPN1.

   host1:pe1(config)#**virtual-router default**
   host1:(config)#**interface tunnel gre:MTI-11 transport-virtual-router pe1**
   host1:(config-if)#**tunnel source 1.1.1.1**
   host1:(config-if)#**tunnel destination 233.3.1.1**
   host1:(config-if)#**tunnel mdt**
   host1:(config-if)#**exit**

   The **tunnel mdt** command enables the IP tunnel component to create an MDT interface.

4. For the VRF, configure an IP interface on the MTI interface.

   host1(config)#**virtual-router pe1**
   host1:pe1(config)#**virtual-router pe1:pe13**
   host1:pe1:pe13(config)#**ip multicast-routing**
   host1:pe1:pe13(config)#**interface loopback 0**
   host1:pe1:pe13(config-if)#**ip address 1.1.1.1 255.255.255.255**
   host1:pe1:pe13(config-if)#**exit**

   NOTE: The IP address that you configure for the VRF must be identical to the IP address of the loopback interface in the parent virtual router.

5. Configure the MTI interface in the VRF.

   host1:pe1:pe13(config)#**interface tunnel gre:MTI-11**
   host1:pe1:pe13(config-if)#**ip unnumbered loopback 0**
   host1:pe1:pe13(config-if)#**ip pim sparse-mode**
   host1:pe1:pe13(config-if)#**exit**

   NOTE: You can use the **ip unnumbered loopback 0** command to configure the MTI interface on the VRF as an unnumbered interface, or the **ip address 1.1.1.1 255.255.255.255** command to configure it as a numbered interface.

6. Specify the GRE tunnel configuration for the parent virtual router.

   NOTE: The MDT interface in the parent virtual router must be unnumbered.

   host1:pe1(config)#**interface tunnel gre:MTI-11.mdt**
   host1:pe1(config-if)#**ip unnumbered loopback 0**
   host1:pe1(config-if)#**ip pim sparse-mode**
   host1:pe1(config-if)#**exit**
   host1:pe1(config)#**virtual-router default**

NOTE: You can configure PIM on IPv4 and IPv6 interfaces. However, IPv6 does not support all PIM configuration options. For information about configuring PIM on IPv6 interfaces, see "Configuring PIM for IPv6 Multicast" on page 273.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Multicast VPNs Overview on page 105
- Configuring Data MDTs on page 109
- Example: Configuring Multicast VPNs on page 111
- Monitoring PIM for IPv4 Multicast on page 121

- **ip pim**
- **tunnel mdt**

Configuring Data MDTs

To configure data MDTs:

1. Configure a dynamic interface profile to specify the PIM configuration of the IP interface or MTI interface in the VRF.

   ```
   host1(config)#profile pe13DataMdtMti
   host1(config-profile)#ip virtual-router pe1:pe13
   host1(config-profile)#ip unnumbered loopback 0
   host1(config-profile)#ip pim sparse-mode
   host1(config-profile)#exit
   ```

2. Configure a dynamic interface profile to specify the IP interface or MDT interface in the parent.

   ```
   host1(config-profile)#profile pe1DataMdtMdt
   host1(config-profile)#ip virtual-router pe1
   host1(config-profile)#ip unnumbered loopback 0
   host1(config-profile)#ip pim sparse-mode
   host1(config-profile)#exit
   ```

3. Configure the destination profile for dynamic IP tunnel creation.

   ```
   host1(config)#gre destination profile pe13DataMdtProfile virtual-router pe1
   host1(config-dest-profile)#tunnel source 1.1.1.1
   host1(config-dest-profile)#tunnel destination ip subnet 233.3.0.0/16
   host1(config-dest-profile)#tunnel mdt profile pe1DataMdtMdt
   host1(config-dest-profile)#profile pe13DataMdtMti
   host1(config-dest-profile)#exit
   ```

   The router uses this destination profile to verify whether it can create a dynamic tunnel, and to supply additional configuration parameters when it creates a tunnel. For more information about creating dynamic IP tunnels, see *Configuring Dynamic IP Tunnels* in the *JunosE IP Services Configuration Guide*.

4. Configure the group address pools in the route map.

```
host1(config)#virtual-router pe1
host1:pe1(config)#ip pim group-address-pool pe13DataMdtGroups 233.3.1.0 233.3.1.255
```

   If the data MDTs are established using ASM, you must divide the range of available MDT P-Group addresses so that PEs source VPN multicasts. All VRFs in a PE draw from a single address pool that contains the range of group addresses assigned to that PE.

   If the data MDTs are established using SSM, you can configure VRFs to transmit on a tunnel using the same MDT P-Group address. Each VRF transmits using a unique P-Source address; however, each data MDT created by the VRF must use a different P-Group address. There might be one sender data MDT and possibly many receiver data MDTs sharing an IP tunnel.

   For SSM, each PE can assign MDT P-Groups from the same range, but the P-Group addresses must be administratively divided among the VPNs as mentioned in the following example:

```
host1(config)#virtual-router pe1
host1:pe1(config)#ip pim group-address-pool pe11DataMdtSSMGroups 233.3.1.0
  233.3.1.255
host1:pe1(config)#ip pim group-address-pool pe12DataMdtSSMGroups 233.3.2.0
  233.3.2.255
host1:pe1(config)#ip pim group-address-pool pe13DataMdtSSMGroups 233.3.3.0
  233.3.3.255
```

5. Configure the access list to match <S,G> and <*,G> entries.

```
host1:pe1(config)#virtual-router pe1:pe13
host1:pe1:pe13(config)#access-list pe13DataMdtSend permit ip host 10.13.0.100
  225.1.1.0 0.0.0.255
host1:pe1:pe13(config)#access-list pe13DataMdt permit ip any 225.1.0.0  0.0.255.255
```

6. Specify a route map to configure the set of <S, G> for which data MDTs can be created, and the threshold to be applied for each SG.

```
host1:pe1:pe13(config)#route-map pe13MdtThresholds permit 10
host1:pe1:pe13(config-route-map)#match ip address pe13DataMdtSend
host1:pe1:pe13(config-route-map)#set threshold 0
host1:pe1:pe13(config-route-map)#route-map pe13MdtThresholds permit 20
host1:pe1:pe13(config-route-map)#match ip address pe13DataMdt
host1:pe1:pe13(config-route-map)#exit
```

7. Configure the data MDT.

```
host1:pe1:pe13(config)#ip pim data-mdt
host1:pe1:pe13(config-ip-pim-data-mdt)#tunnel source 1.1.1.1
host1:pe1:pe13(config-ip-pim-data-mdt)#tunnel group-address-pool  pe13DataMdtG$
host1:pe1:pe13(config-ip-pim-data-mdt)#route-map pe13MdtThresholds
```

NOTE: You can configure PIM on IPv4 and IPv6 interfaces. However, IPv6 does not support all PIM configuration options. For information about configuring PIM on IPv6 interfaces, see "Configuring PIM for IPv6 Multicast" on page 273.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Multicast VPNs Overview on page 105
- Configuring the Default MDT on page 107
- Example: Configuring Multicast VPNs on page 111
- Monitoring PIM for IPv4 Multicast on page 121
- **ip pim data-mdt**
- **ip pim group-address-pool**
- **mdt-data-delay**
- **mdt-data-holddown**
- **mdt-data-timeout**
- **mdt-interval**
- **set threshold**
- **tunnel group-address-pool**

Example: Configuring Multicast VPNs

The following example illustrates how to configure multicast VPNs.

- Requirements on page 111
- Overview on page 111
- Configuring a Multicast VPN Network on page 112

Requirements

This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

In the following example (Figure 13 on page 112), customer edge router 1 (CE1) and customer edge router 2 (CE2) exist in two separate VPNs. Each VPN is configured with its assigned Multicast Domain (235.1.1.1 and 235.1.1.2, respectively).

Figure 13: Multicast VPNs

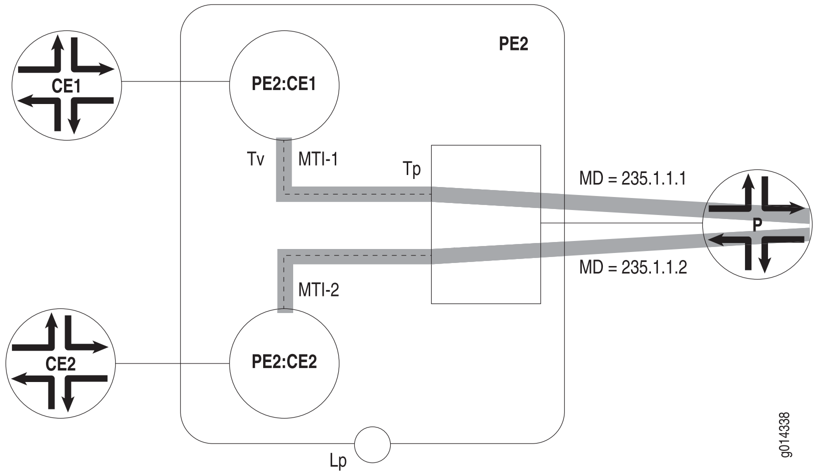

To better understand the example, keep the following in mind:

- Lp is a loopback interface in the parent router. This address is the loopback interface used as the BGP peer address of the provider edge router (PE). Its address is advertised in the provider address space.
- Tv is the MTI in the VRF. This interface is typically configured as a PIM sparse-mode interface (though you can configure it for dense-mode or sparse-dense-mode). Any packets that originate in the VRF are sent using the address of this interface as the source address. You must set this interface address to be identical to loopback interface of the parent router (Lp).

CAUTION: **Defining the Tv interface with an address other than the loopback interface of the parent router might restrict operation with non-Juniper Networks routers.**

- Tp is an unnumbered IP interface that is tied to the loopback interface of the provider edge router (PE).

Configuring a Multicast VPN Network

Step-by-Step Procedure

To configure the example, use the following general procedures:

NOTE: This example provides general information for configuring a simple Multicast VPN network. For detailed information about creating GRE tunnels, see *Configuring IP Tunnels* in the *JunosE IP Services Configuration Guide*. For detailed information about PIM sparse-mode configuration, see *PIM Sparse Mode* in “Understanding PIM for IPv4 Multicast” on page 90.

1. Configure BGP/MPLS VPN.


```
host1:PE2(config-router)#router bgp 100
host1:PE2(config-router)#address-family vpnv4 unicast
host1:PE2(config-router-af)#neighbor 1.1.1.1 activate
host1:PE2(config-router-af)#neighbor 1.1.1.1 next-hop-self
host1:PE2(config-router-af)#neighbor 3.3.3.3 activate
host1:PE2(config-router-af)#neighbor 3.3.3.3 next-hop-self
host1:PE2(config-router-af)#exit-address-family
```

See *Configuring BGP-MPLS Applications* in the *JunosE BGP and MPLS Configuration Guide* for details.

2. Configure PIM sparse mode in the core and RP for MVPN group addresses.

NOTE: For MVPN, it is a typical practice to use shared trees.

```
host1:PE1(config-router)#virtual-router PE2
host1:PE2(config)#ip multicast-routing
host1:PE2(config)#
host1:PE2(config)#! MDT RP is 72.72.72.72 (P1)
host1:PE2(config)#access-list 1 permit ip 235.0.0.0 0.255.255.255 any
host1:PE2(config)#ip pim rp-address 72.72.72.72 1
host1:PE2(config)#
host1:PE2(config)#! Do not switch from RPT for MDTs
host1:PE2(config)#ip pim spt-threshold infinity group-list 1
host1:PE2(config)#
```

3. Configure the loopback interface, Lp, in parent router PE2.

```
host1:PE2(config)#interface loopback 0
host1:PE2(config-if)#ip address 2.2.2.2 255.255.255.255
host1:PE2(config-if)#ip pim sparse-mode
host1:PE2(config-if)
```

NOTE: You must configure the loopback interface for PIM sparse mode to support unnumbered MDTs.

4. Add PIM-SM to core-facing interfaces.

```
host1:PE2(config)#interface atm2/1.20
host1:PE2(config-subif)#ip pim sparse-mode
host1:PE2(config-subif)#
```

5. Extend the BGP router configuration to contribute VPN routes into the multicast router table of the VRF using the **ip route-type both** command.

```
host1:PE2(config)#router bgp 100
host1:PE2(config-router)#address-family ipv4 unicast vrf PE21
host1:PE2(config-router-af)#ip route-type both
host1:PE2(config-router-af)#exit
host1:PE2(config-router)#
```

6. Configure the GRE tunnel for VPN1.

```
host1(config)#interface tunnel gre:MTI-21 transport-virtual-router PE2
host1(config-if)#tunnel source 2.2.2.2
host1(config-if)#tunnel destination 235.1.1.1
host1(config-if)#tunnel mdt
host1(config-if)#exit
host1(config)#
```

7. Configure the GRE tunnel for VPN2.

```
host1(config)#interface tunnel gre:MTI-22 transport-virtual-router PE2
host1(config-if)#tunnel source 2.2.2.2
host1(config-if)#tunnel destination 235.1.1.2
host1(config-if)#tunnel mdt
host1(config-if)#exit
host1(config)#
```

8. Configure the IP interface (Tv) in PE2:CE1 as a numbered or unnumbered PIM sparse-mode interface. Use the same address as the loopback interface, Lp in the parent router, PE2.

```
host1(config)#virtual-router PE2:CE1
host1:PE2:CE1(config)#interface tunnel gre:MTI-21
host1:PE2:CE1(config-if)#ip address 2.2.2.2 255.255.255.255
host1:PE2:CE1(config-if)#ip pim sparse-mode
host1:PE2:CE1(config-if)#exit
host1:PE2:CE1#
```

9. Configure the IP interface (Tv) in PE2:CE2 as a numbered or unnumbered PIM sparse-mode interface. Use the same address as the loopback interface, Lp in the parent router, PE2.

```
host1(config)#virtual-router PE2:CE2
host1:PE2:CE2(config)#interface loopback 0
host1:PE2:CE2(config-if)#ip address 2.2.2.2 255.255.255.255
host1:PE2:CE2(config-if)#exit
host1:PE2:CE2(config)#
host1:PE2:CE2(config)#interface tunnel gre:MTI-22
host1:PE2:CE2(config-if)#ip unnumbered loopback 0
host1:PE2:CE2(config-if)#ip pim sparse-mode
host1:PE2:CE2(config-if)#exit
host1:PE2:CE2#
```

10. Configure the Tp interfaces in the parent router, PE2, as unnumbered PIM sparse-mode interfaces tied to the loopback interface, Lp.

```
host1(config)#virtual-router PE2
host1: PE2(config)#interface tunnel gre:MTI-21.mdt
host1:PE2(config-if)#ip unnumbered loopback 0
host1:PE2(config-if)#ip pim sparse-mode
host1:PE2(config-if)#exit
host1:PE2(config)#

host1:PE2(config)#interface tunnel gre:MTI-22.mdt
host1:PE2(config-if)#ip unnumbered loopback 0
host1:PE2(config-if)#ip pim sparse-mode
host1:PE2(config-if)#exit
host1:PE2(config)#
```


NOTE: You can configure PIM on IPv4 and IPv6 interfaces. However, IPv6 does not support all PIM configuration options. For information about configuring PIM on IPv6 interfaces, see "Configuring PIM for IPv6 Multicast" on page 273.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Multicast VPNs Overview on page 105
- Configuring the Default MDT on page 107
- Configuring Data MDTs on page 109
- Monitoring PIM for IPv4 Multicast on page 121

- **tunnel mdt**

Configuring PIM Sparse Mode Join Filters for IPv4

You can use PIM sparse mode join filters to prevent multicast state from being created in the PIM sparse mode router. The filters are applied to join entries in PIM join/prune messages that are received from PIM sparse mode neighbors.

By denying joins at the edge of a network, you can limit the multicast state and traffic in the network. By accepting only certain joins, you can control which multicast services an end user can receive. PIM join filters also reduce the potential for denial of service (DoS) attacks where large numbers of joins forwarded to each router on the RPT can result in a PIM state explosion and very high memory consumption.

For information about how to create access lists, see *Configuring Routing Policy* in the *JunosE IP Services Configuration Guide*.

To specify an extended access list that you want this PIM interface to use as a join filter:

- Issue the **ip pim join-filter** command in Global Configuration mode:

  ```
  host1(config)#ip pim join-filter gold
  ```

  The **no** version removes the filter association. You can apply the join filter at the global level or at the interface level. If an interface-level filter exists, it takes precedence over the global-level filter.

NOTE: You can configure PIM on IPv4 and IPv6 interfaces. However, IPv6 does not support all PIM configuration options. For information about configuring PIM on IPv6 interfaces, see "Configuring PIM for IPv6 Multicast" on page 273.

Related Documentation

- **ip pim join-filter**

Configuring PIM for IPv4 SSM

Source Specific Multicast (SSM) is a datagram delivery model that best supports one-to-many applications, also known as broadcast applications. SSM is networking technology that targets audio and video broadcast application environments.

To configure PIM SSM, you enable PIM SSM on the router and define the SSM range of IP multicast addresses.

To use PIM SSM, IGMPv3 must be configured on customer premise equipment (CPE)–facing interfaces to receivers, and PIM sparse mode must be configured on CPE-facing interfaces to sources and on core-facing interfaces. After configuring SSM, you can use the **show ip pim sparse-mode sg-state** command to display SSM group membership information.

To configure PIM SSM:

1. Enable PIM SSM on the E Series router. The IANA SSM range (232.0.0.0/8) is configured by default. You can modify the SSM address range by using the access list.

   ```
   host1(config)#access-list 15 permit ip any host 239.0.0.2
   host1(config)#access-list 15 permit ip any 232.0.0.0 0.225.225.225
   host1(config)#ip pim ssm range 15
   ```
2. Enable PIM sparse mode on the CPE-facing interface towards the source or core.
3. Enable IGMPv3 on the CPE-facing interface towards the receiver. PIM SSM also works with IGMPv2 if you configure the ssm-map in IGMP as in the following example:

PIM SSM also works with IGMPv2 if you configure the ssm-map in IGMP as in the following example:

```
host1(config)#ip pim ssm
host1(config)#access-list ssm_map1 permit 232.0.0.1 255.255.255.255
host1(config)#ip igmp ssm-map enable
host1(config)#ip igmp ssm-map static ssm_map1 51.0.0.1
```

The **no** version disables ssm-map:

host1(config)#**no ip igmp ssm-map static ssm_map1 51.0.0.1**

NOTE: You can configure PIM on IPv4 and IPv6 interfaces. However, IPv6 does not support all PIM configuration options. For information about configuring PIM on IPv6 interfaces, see "Configuring PIM for IPv6 Multicast" on page 273.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Monitoring PIM for IPv4 Multicast on page 121
- **ip pim ssm**

BFD Protocol for PIM for IPv4 Overview

The **ip pim bfd-liveness-detection** command configures the Bidirectional Forwarding Detection (BFD) protocol for PIM. The BFD protocol uses control packets and shorter detection time limits to more rapidly detect failures in a network. Also, because they are adjustable, you can modify the BFD timers for more or less aggressive failure detection.

PIM routers send periodic hello messages from each PIM-enabled interface. You can configure this interval using the **ip pim query-interval** command. By default, the PIM router sends a hello message every 30 seconds (with an interval range of 0–210 seconds). If it receives no response from a neighbor within 3.5 times the interval value (a minimum of 3.5 seconds), the PIM router drops the neighbor.

In contrast, when a BFD session exists between neighbors, a PIM neighbor that goes down is detected quickly (in milliseconds rather than in seconds).

When you issue the **ip pim bfd-liveness-detection** command on a PIM router, the router establishes BFD liveness detection with all BFD-enabled PIM neighbors. When the local router receives an update from a remote PIM neighbor—if BFD is enabled and if the session is not already present—the local router attempts to create a BFD session to the remote neighbor.

Each adjacent pair of neighbors negotiates an acceptable transmit interval for BFD packets. The negotiated value can be different on each neighbor. Each neighbor then calculates a BFD liveness detection interval. When a neighbor does not receive a BFD packet within the detection interval, it declares the BFD session to be down.

NOTE: Before the router can use the **ip pim bfd-liveness-detection** command, you must specify a BFD license key. To view an already configured license, use the **show license bfd** command.

For general information about configuring and monitoring the BFD protocol, see *Configuring BFD* in the *JunosE IP Services Configuration Guide*.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Configuring the BFD Protocol for PIM for IPv4 on page 118
- Monitoring PIM for IPv4 Multicast on page 121
- **ip pim bfd-liveness-detection**

Configuring the BFD Protocol for PIM for IPv4

To enable BFD (bidirectional forwarding detection) and define BFD values to more quickly detect PIM data path failures:

- Issue the **ip pim bfd-liveness-detection** command in Interface Configuration mode.

  ```
  host1(config-if)#ip pim bfd-liveness-detection minimum-interval 800
  ```

NOTE: Before the router can use the **ip pim bfd-liveness-detection** command, you must specify a BFD license key. To view an already configured license, use the **show license bfd** command.

For details on liveness detection negotiation, see *Negotiation of the BFD Liveness Detection Interval* in the *JunosE IP Services Configuration Guide*.

NOTE: You can configure PIM on IPv4 and IPv6 interfaces. However, IPv6 does not support all PIM configuration options. For information about configuring PIM on IPv6 interfaces, see "Configuring PIM for IPv6 Multicast" on page 273.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- BFD Protocol for PIM for IPv4 Overview on page 117
- Monitoring PIM for IPv4 Multicast on page 121
- **ip pim bfd-liveness-detection**

Removing PIM for IPv4

To remove PIM from the VR:

- Issue the **no router pim** command in Global Configuration mode.

  ```
  host1:boston(config)#no router pim
  ```

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Enabling PIM for IPv4 on a Virtual Router on page 97
- Disabling PIM for IPv4 on a Virtual Router on page 98

- Enabling PIM for IPv4 on an Interface on page 98
- Resetting PIM Counters and Mappings for IPv4 on page 119
- Monitoring PIM for IPv4 Multicast on page 121
- **router pim**

Resetting PIM Counters and Mappings for IPv4

You can use the **clear ip pim** commands to reset PIM counters and mappings.

- Issue the **clear ip pim auto-rp** command in Privileged Exec mode to clear the group-to-RP router mappings that the router learned through auto-RP:

  host1#**clear ip pim auto-rp 192.34.56.7**

  There is no **no** version.

  NOTE: Specify the IP address of an RP to clear the group-to-RP mappings for a particular RP. If you do not specify an IP address, the router clears the group-to-RP mappings on all RP routers learned through auto-RP.

- Issue the **clear ip pim remote-neighbor count** command in Privileged Exec mode to clear the counters for remote neighbor statistics on all interfaces or the specified interface:

  host1#**clear ip pim interface atm 3/0.5 count**

  There is no **no** version.

- Issue the **clear ip pim interface count** command in Privileged Exec mode to clear the counters for multicast packet statistics on all interfaces or a specified interface:

  host1#**clear ip pim interface atm 3/0.5 count**

  There is no **no** version.

  NOTE: Specify the IP address of an interface to clear the counters for that interface. If you do not specify an interface, the router clears the counters on all interfaces.

Related Documentation

- Understanding PIM for IPv4 Multicast on page 90
- Enabling PIM for IPv4 on a Virtual Router on page 97
- Disabling PIM for IPv4 on a Virtual Router on page 98
- Enabling PIM for IPv4 on an Interface on page 98
- Removing PIM for IPv4 on page 118
- Monitoring PIM for IPv4 Multicast on page 121

- **router pim**

CHAPTER 6

Monitoring PIM for IPv4 Multicast

You can display information about PIM events and parameters. You can use the debug PIM commands to view information about PIM events and the **show ip pim** commands to display information about PIM settings.

Enabling the Display of a PIM Event

To display information about the selected event.

- Issue the **debug ip pim** command in Privileged Exec mode.

  host1#**debug ip pim events severity 1 verbosity low**

  Use the **no** version to disable the display.

NOTE: To control the type of events displayed, specify a severity level, and to control how much information to display, specify a verbosity level.

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **debug ip pim**

Disabling the Display of a PIM Event

To turn off the display of information previously enabled with the **debug ip pim** command.

- Issue the **undebug ip pim** command in the Privileged Exec mode.

  host1#**undebug ip pim events**

  There is no **no** version to disable the display.

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **undebug ip pim**

Monitoring PIM Router-Level Information for IPv4

Purpose Display general PIM router-level information.

Action To display general PIM router-level information:

```
host1:1#show ip pim
Default PIM Version: 2
Default Domain Id: 0
Default Hello Period: 30
Default Hello HoldTime: 105
Join-Prune Interval: 100
Join-Prune Holdtime: 350
Keepalive Period: 210
Assert Time: 210
Register Suppression Time: 60
Register Probe Time: 5
Register TTL: 64
SSM enabled, range default
Sparse-Mode Graceful Restart Duration: 30
Graceful restart is complete (timer 0 seconds)
Join filter, access-list bronze
Designated Router Priority: 2
```

Meaning Table 29 on page 122 lists the **show ip pim** command output fields.

Table 29: show ip pim Output Fields

Field Name	Field Description
Default PIM Version	Default PIM version number (always 2)
Default Domain Id	Default Domain Id (always 0)
Default Hello period	Default interval (in minutes) at which the router sends hello messages to neighbors

Table 29: show ip pim Output Fields *(continued)*

Field Name	Field Description
Default Hello Hold Time	Default time (in minutes) for which the router keeps the neighbor state alive
Join-Prune Interval	Interval value (in seconds) set in the join/prune message originated by the PIM router
Join-Prune Holdtime	Hold time value (in seconds) set in the join/prune messages originated by the PIM router. The hold time is 3.5 times the PIM join/prune message interval value.
Keepalive Period	Time SG join state is maintained in the absence of SG Join message
Assert Time	Period after last assert before assert state is timed out
Register Suppression Time	Period during which a designated router stops sending registers to the RP
Register Probe Time	Time before register suppression time (RST) expires when a designated router might send a NULL-Register to the RP
Register TTL	TTL value (in PIM register packets) originated by this PIM router
SSM	State of SSM on this PIM router (enabled or disabled)
range	Default SSM group range or name of the access list specifying the range
Sparse-Mode Graceful Restart Duration	Restart interval in seconds
Join filter	Name of the join filter access-list (if configured) for this PIM router
Designated Router Priority	Designated router priority value

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim**

Monitoring RP Routers and the RP Mapping Agent in a PIM Sparse Mode Environment for IPv4

Purpose Display information about RP routers and the RP mapping agent in a PIM sparse mode environment.

Action To display information about RP routers and the RP mapping agent in a PIM sparse mode environment.

To display information on a PIM router that is an Auto RP mapping agent:

```
host1:1#show ip pim auto-rp
This PIM router is an Auto RP mapping agent.
  Configured with ttl 64
  [ Using interface addr 121.0.0.1, interval 60 ].
PIM AutoRP candidate RP mapping(s)
```

To display information on a PIM router that is not an Auto RP mapping agent:

```
host1:1# show ip pim auto-rp
This PIM router is _not_ an Auto RP mapping agent.
PIM AutoRP candidate RP mapping(s)
Candidate RP 122.0.0.1
  Group(s) 224.0.0.0/4, AutoRP, ttl 64, interval 60, from access List 1
Candidate RP 122.0.0.1
  Group(s) 224.0.1.39/32 (negative), AutoRP, ttl 64, interval 60, from access
List 1
Candidate RP 122.0.0.1
  Group(s) 224.0.1.40/32 (negative), AutoRP, ttl 64, interval 60, from access
List 1
```

Meaning Table 30 on page 124 lists the **show ip pim auto-rp** command output fields.

Table 30: show ip pim auto-rp Output Fields

Field Name	Field Description
Configured with ttl	Number of hops for which the RP discovery message is valid
Using interface addr	IP address of the interface from which the router sends RP discovery messages
interval	Time interval, in seconds, at which the router sends RP discovery messages
PIM AutoRP candidate RP mapping(s)	Routers that the RP mapping agent is evaluating to determine an RP router for this interface

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim auto-rp**

Monitoring BSR Information for IPv4

Purpose Display BSR information and the group prefixes for which the local router is a candidate RP in a PIM sparse mode environment.

Action To display information on a router that is the elected BSR:


```
host1:1#show ip pim bsr

This PIM router is a Candidate BSR.
  Configured on intf ATM3/0.101, address: 101.0.0.1
  hashMaskLen 30, priority 2, period 60 seconds.
Elected BSR is this router, next BSM in 3 seconds.

Local candidate RP mapping(s):
Candidate RP 101.0.0.1
  224.0.0.0/4, BSR, hold-time 150, interval 60, priority 192
  228.0.0.0/24, BSR, hold-time 150, interval 60, priority 192, from access-list
acl
  230.0.0.0/24, BSR, hold-time 150, interval 60, priority 192, from access-list
acl
```

To display information on a router that is a candidate BSR:

```
host1:1#show ip pim bsr
This PIM router is a Candidate BSR.
  Configured on intf ATM3/0.100, address: 100.0.0.1
  hashMaskLen 30, priority 2, period 60 seconds.
Elected BSR is 101.0.0.1 (priority 0), expires in 73 seconds.
```

To display information on a router that is not a candidate BSR:

```
host1:1#show ip pim bsr
This PIM router is not a Candidate BSR.
Elected BSR is 101.0.0.1 (priority 0), expires in 73 seconds.
```

Meaning Table 31 on page 125 lists the **show ip pim bsr** command output fields.

Table 31: show ip pim bsr Output Fields

Field Name	Field Description
Candidacy	Whether the router is a candidate BSR
Configured on	Interface on which the router is configured
address	Address of the router
hashMaskLen	Hash mask length
priority	Priority of the router
period	Time between bootstrap messages, in seconds
Elected BSR	This router or IP address of the elected bootstrap router
next BSM	If BSR is this router, time until the next bootstrap message is sent, in seconds
expires in	If BSR is not this router, time until the elected BSR expires if no bootstrap messages are received
Local candidate RP mapping(s)	Routers that the mapping agent is evaluating to determine an RP router for this interface

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim bsr**

Monitoring Active Data MDTs for IPv4

Purpose Display information about active data MDTs. Use the following keywords to display the required information:

- **sender**—To display information about data MDTs on which the provider edge transmits data
- **receiver**—To display information about data MDTs on which the provider edge receives data
- **group**—To display information about an IP PIM group address pool
- **summary**—To display a summary of configuration for each data MDT
- **count**—To display the number of data MDTs

Action To display information about a data MDT sender:

```
host1:PE1#show ip pim data-mdt 225.1.1.1
PE11 - Sender
  C-SG: 10.11.0.100, 225.1.1.1
  P-SG: 1.1.1.1, 235.0.1.1
  MTI: TUNNEL gre:mvpn-dynamic-1
    Data rate/Threshold: 10012/500 Kbps
    Time until next MDT Join TLV: 25 seconds
```

To display information about a data MDT receiver:

```
host1:PE1#show ip pim data-mdt 225.2.2.2
PE31 - Receiver
  C-SG: 10.13.0.100, 225.2.2.2
  P-SG: 3.3.3.3, 235.0.1.1
  MTI: TUNNEL gre:mvpn-dynamic-3
    Time until MDT Join TLV expires: 29 seconds
```

To display a summary of data MDT senders:

```
host1:PE1#show ip pim data-mdt senders summary

VRF    S/R       C-Group     C-Source      P-Group    P-Source          MTI

----   -------   ---------   -----------   ---------  --------   ---------------
PE11   Sender   225.1.1.1   10.11.0.100   235.0.1.1  1.1.1.1    TUNNEL
gre:mvpn-dynamic-1
PE12   Sender   225.1.1.1   10.12.0.100   235.0.1.2  1.1.1.1    TUNNEL
gre:mvpn-dynamic-2

Counts: 2 senders, 0 receivers, total 3.
```

To display the number of data MDT senders and receivers:

```
host1:PE1#show ip pim data-mdt count
Counts: 2 senders, 1 receivers, total 3.
```

Meaning Table 32 on page 127 lists the **show ip pim data-mdt** command output fields.

Table 32: show ip pim data-mdt Output Fields

Field Name	Field Description
PE *Name*	Name of the PE
C-SG	Address of the C-SG
P-SG	Address of the P-SG
MTI	Name of the dynamic IP tunnel on which the data MDT was created
Data rate/Threshold	Rate and threshold of multicast data
Time until next MDT Join TLV	Configured delay until next MDT Join TLV
Time until MDT Join TLV expires	Configured delay until MDT Join TLV expires
Time until switchover from default-MDT	Configured delay until the data MDT switches over to the default MDT

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim data-mdt**

Monitoring Each (Source, Group) Pair for PIM Dense Mode for IPv4

Purpose Display information for each (Source, Group) pair for PIM dense mode.

Action To display information for each (Source, Group) pair for PIM dense mode:

```
host1:8#show ip pim dense-mode sg-state
PIM DM route table and pruned oif information
<122.0.0.1, 224.0.1.39>   EntryExpires: 99
  RPF Route: 122.0.0.0/255.0.0.0   IIF: 107.0.8.4   UpNbr: 107.0.4.8
  Pruned Oifs:
   Address: 108.0.8.5   IfId: 95
    Pruned due to assert
    Pruned time remaining 129
<130.0.0.2, 224.0.1.39>   EntryExpires: 100
  RPF Route: 130.0.0.0/255.0.0.0   IIF: 107.0.8.4   UpNbr: 107.0.4.8
  Pruned Oifs:
   Address: 108.0.8.5   IfId: 95
    Pruned due to assert
    Pruned time remaining 130
<121.0.0.1, 224.0.1.40>   EntryExpires: 102
  RPF Route: 121.0.0.0/255.0.0.0   IIF: 107.0.8.4   UpNbr: 107.0.4.8
  Pruned Oifs:
   Address: 108.0.8.5   IfId: 95
    Pruned due to assert
    Pruned time remaining 133
```

Meaning Table 33 on page 128 lists the **show ip pim dense-mode sg-state** command output fields.

Table 33: show ip pim dense-mode sg-state Output Fields

Field Name	Field Description
(Source, Group) pair	IP address and network mask for the unicast route
EntryExpires	Time until the (S,G) pair entry expires
RPF Route	Reverse-path forwarding route
IIF	IP address of incoming interface
UpNbr	IP address of upstream neighbor
Pruned Oifs	Outgoing interfaces that have been pruned
Address	IP address of outgoing interface
IfId	Index of the interface
Pruned due to	Reason for prune: assert or explicit prune
Pruned time remaining	Time in seconds until the prune expires

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim dense-mode sg-state**.

Monitoring PIM Interfaces for IPv4

Purpose

Display information about PIM interfaces. Specify no keywords or variables to view information about all PIM interfaces.

Specify the **summary** keyword to view the number of configured, enabled, and disabled PIM dense mode, PIM sparse mode, and PIM sparse-dense mode interfaces.

Specify the **count** keyword to view the number of multicast packets that the interface has sent and received.

Action

To display information about PIM interfaces. Specify no keywords or variables to view information about all PIM interfaces:

```
host1# show ip pim interface
PIM Interface Table
Interface Addr Interface   State Ver Mode    Nbr    Hello  J/P    DR    DR Addr
  JoinFilter
                                             Count  Intvl  Intvl  Pri
101.0.0.1      ATM2/0.100  Up      2  Sparse  1      30     150    5    101.0.0.2
  -
102.0.0.1      ATM2/0.101  Up      2  Sparse  1      30     100    2    102.0.0.2
  silver
```

```
103.0.0.1      ATM3/0.102  Up     2  Sparse   1      30    100     2    103.0.0.1
  gold

host1#show ip pim interface summary
PIM Interface Summary
  SM:     0, 0 enabled, 0 disabled
  DM:     0, 0 enabled, 0 disabled
  SM/DM: 1, 0 enabled, 1 disabled

host1#show ip pim interface count
PIM Interface Count
Interface Addr   Interface Name       ControlPktCount In|Out
                                      Hello     JoinPrune  Assert
192.32.10.20     ATM3/0.20             0         0           0
                                       0         0           0
```

Meaning Table 34 on page 129 lists the **show ip pim interface** command output fields.

Table 34: show ip pim interface Output Fields

Field Name	Field Description
Interface Addr	IP address of the interface
Interface Name	Type and specifier of the interface. For details about interface types and specifiers, see Interface Types and Specifiers in JunosE Command Reference Guide.
Ver	Version of PIM running on this interface
Mode	PIM mode running on this interface: Sparse, Dense, or SparseDense
Nbr Count	Number of neighbors connected to this interface
Hello Intvl	Time interval, in seconds, at which the interface sends hello messages to neighbors
J/P Interval	Time interval, in seconds, at which the interface sends the join/prune messages
DR Addr	Address of the designated router
SM	Number of PIM sparse mode interfaces
DM	Number of PIM dense mode interfaces
SM/DM	Number of PIM sparse-dense mode interfaces: • enabled—Number of interfaces administratively enabled • disabled—Number of interfaces administratively disabled
DR Pr	Designated router priority value

Table 34: show ip pim interface Output Fields *(continued)*

Field Name	Field Description
ControlPktCount In|Out	PIM messages received on and sent from this interface
Hello	Total number of hello messages
JoinPrune	Total number of join and prune messages
Assert	Total number of assert messages

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim interface**

Monitoring PIM Neighbors for IPv4

Purpose Display information about PIM neighbors that the router discovered.

Action To display information about PIM neighbors that the router discovered:

host1#**show ip pim neighbor**

```
Neighbor
  Addr      Interface Name    Uptime    Expires    Ver    Mode    BFD
--------    ---------------  --------   --------   ---   ------   ---
1.1.1.1     FastEthernet1/1  00:02:49   00:01:27   2     Sparse
Neighbor
  Addr      Interface Name    Uptime    Expires    Ver    Mode    BFD
--------    ---------------  --------   --------   ---   ------   ---
1.1.1.1     FastEthernet1/1  00:03:16   00:01:30   2     Sparse   up
Neighbor
  Addr      Interface Name    Uptime    Expires    Ver    Mode    BFD
--------    ---------------  --------   --------   ---   ------   ----
1.1.1.1     FastEthernet1/1  00:00:07   00:01:39   2     Sparse   down
```

Meaning Table 35 on page 130 lists the **show ip pim neighbor** command output fields.

Table 35: show ip pim neighbor Output Fields

Field Name	Field Description
Neighbor Addr	IP address of the neighbor
Interface Name	Type and specifier of the interface to which the neighbor connects. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*
Uptime	Time since the router discovered this neighbor in *days hours:minutes:seconds* format

Table 35: show ip pim neighbor Output Fields *(continued)*

Field Name	Field Description
Expires	Time available for the neighbor to send a hello message to the interface. If the neighbor does not send a hello message during this time, it no longer is a neighbor
Ver	Version of PIM that the neighbor is running
Mode	PIM mode that the neighbor is using: Sparse, Dense, or SparseDense
BFD	BFD status: up or down

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim neighbor**

Monitoring PIM Group-to-RP Mappings for IPv4

Purpose Display information about PIM group-to-RP mappings. Specify the address of a group to view PIM group-to-RP mappings for a particular group. Specify the mapping keyword to display all group-to-RP mappings that the router has recorded.

Action To display information about PIM group-to-RP mappings:

```
host1:8#show ip pim rp mapping
PIM Group-to-RP mapping(s)
Group(s) 224.0.0.0/4
  RP 122.0.0.1,  priority 0,  via  AutoRP, expiryTime 88
Group(s) 224.0.1.39/32 (negative)
  RP 122.0.0.1,  priority 0,  via  AutoRP (Negative), expiryTime 88
Group(s) 224.0.1.40/32 (negative)
  RP 122.0.0.1,  priority 0,  via  AutoRP (Negative), expiryTime 88

host1:8#show ip pim rp mapping
PIM Group-to-RP mapping(s)
Group(s) 224.0.0.0/4
  RP 134.0.0.1, priority 0, via Static, from access-list 1
Group(s) 232.0.0.0/16
  RP 134.0.0.1, priority 0, via BSR, expires in 121 seconds
```

Meaning Table 36 on page 131 lists the **show ip pim rp** command output fields.

Table 36: show ip pim rp Output Fields

Field Name	Field Description
Group(s)	Prefix of the multicast group
RP	IP address of RP router for the multicast group

Table 36: show ip pim rp Output Fields *(continued)*

Field Name	Field Description
priority	Priority of the router
via	Method by which the RP router was assigned: AutoRP, Static RP, or BSR
expiryTime	Time in seconds at which the RP mapping becomes invalid, unless the mapping agent (access list) reassigns the RP router to this group

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim rp**

Monitoring the RP Router that a Multicast Group Uses for IPv4

Purpose Show which RP router that a multicast group is using.

Action To show which RP router that a multicast group is using:

```
host1:2#show ip pim rp-hash 232.1.1.1
Group(s) 224.0.0.0/4
  RP 122.0.0.1,  priority 0,  via  AutoRP, expiryTime 128

host1:2#show ip pim rp-hash 226.0.0.1
Group(s) 226.0.0.0/24
  RP 101.0.0.1, priority 192, via BSR, expires in 145 seconds
 *RP 145.0.0.3, priority 192, via BSR, expires in 145 seconds
```

Meaning Table 37 on page 132 lists the **show ip pim rp-hash** command output fields.

Table 37: show ip pim rp-hash Output Fields

Field Name	Field Description
Group(s)	Multicast group or groups
RP	RP router for the multicast group
priority	Priority of the router
via	Method by which the RP router was assigned: AutoRP, Static RP, or BSR
expiryTime	Time in seconds at which the RP mapping becomes invalid, unless it is renewed by the mapping agent

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim rp-hash**

Monitoring each (S,G) Pair for PIM Sparse Mode and PIM SSM for IPv4

Purpose Display information for each (S,G) pair for PIM sparse mode and PIM SSM.

Action To display information for each (S,G) pair for PIM sparse mode and PIM SSM:

```
host1:2# show ip pim sparse-mode sg-state
<*, 224.0.1.40>
  Group-to-RP mapping: 224.0.0.0/240.0.0.0   RP: 123.0.0.1
  RPF Route: 123.0.0.0/255.0.0.0   IIF: 106.0.7.3   UpNbr: 106.0.3.7
  Oifs:
   Auto RP Discovery SELF oif.
    Joined as <*, G>
<*, 225.1.2.3>
  Group-to-RP mapping: 224.0.0.0/240.0.0.0   RP: 123.0.0.1
  RPF Route: 123.0.0.0/255.0.0.0   IIF: 106.0.7.3   UpNbr: 106.0.3.7
  Oifs:
   Address: 78.7.7.7   Interface: loopback7
    Local group membership present.
<*, 235.1.1.1>
  Group-to-RP mapping: 224.0.0.0/240.0.0.0   RP: 123.0.0.1
  RPF Route: 123.0.0.0/255.0.0.0   IIF: 106.0.7.3   UpNbr: 106.0.3.7
  Oifs:
   Address: 78.7.7.7   Interface: loopback7
    Local group membership present.
<118.1.33.34, 232.0.0.1>
  SSM Group
  RPF Route: 118.1.0.0/255.255.0.0   IIF: 118.1.0.1 (Directly attached)
  Oifs:
   Register Oif to RP: 141.0.0.2 suppressed for SSM Group.
   Address: 134.0.0.1   Interface: ATM3/0.104
    Joined as <S, G>   Join Expires: 161
<118.1.33.35, 232.0.0.1>
  SSM Group
  RPF Route: 118.1.0.0/255.255.0.0   IIF: 118.1.0.1 (Directly attached)
  Oifs:
   Register Oif to RP: 141.0.0.2 suppressed for SSM Group.
   Address: 134.0.0.1   Interface: ATM3/0.104
    Joined as <S, G>   Join Expires: 161
<10.0.1.8, 235.1.1.1>      EntryExpires: 143
  Group-to-RP mapping: 224.0.0.0/240.0.0.0   RP: 123.0.0.1
  RPF Route: 10.0.0.0/255.0.0.0   IIF: 106.0.7.3   UpNbr: 106.0.3.7
  Oifs:
   Address: 78.7.7.7   Interface: loopback7
    Joined as <*, G>
Count of entries - <S, G>    : 3
                   <*, G>    : 3
                   <*, *, RP>: 0
```

Meaning Table 38 on page 134 lists the **show ip pim sparse-mode sg-state** command output fields.

Table 38: show ip pim sparse-mode sg-state Output Fields

Field Name	Field Description
(S, G) pair	Source, Group pair for which information is provided
Group-to-RP mapping	IP addresses and network mask of the multicast group
RP	IP address of RP router
SSM group	Indicator that this is an SSM group
RPF Route	IP address and network mask of the RPF route
IIF	IP address of the incoming interface for the RPF route
UpNbr	IP address of the upstream neighbor
Oifs	Outgoing interface
Auto RP Discovery SELF oif	Indicates that RP router for this group was assigned through auto-RP
Register Oif to RP	IP address of RP router for the outgoing interface; suppressed for SSM
Address	IP address of outgoing interface
Interface	Type and specifier of the interface. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*
Joined as	Type of mapping: • (S, G)—Mapping from a specific source to a specific group • (*, G)—Mapping from any source to a specific group • (*, * , RP)—Mapping from any source to any group
Join expires	Number of seconds before the (S,G) membership expires
Count of entries	Total counts of (S,G) pair mappings

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim sparse-mode sg-state**

Monitoring Unicast Routes that PIM Sparse Mode Uses for IPv4

Purpose Display the unicast routes that PIM sparse mode is using.

Action To display the unicast routes that PIM sparse mode is using:

```
host1:2#show ip pim sparse-mode unicast-route
PIM SM unicast route table information
Route                          RpfNbr          Iif              Pref  Metric
----------------------------------------------------------------------------
122.0.0.0      /255.0.0.0                       122.0.0.1       255   1
Count of entries: 1
```

Meaning Table 39 on page 135 lists the **show ip pim sparse-mode unicast-route** command output fields.

Table 39: show ip pim sparse-mode unicast-route Output Fields

Field Name	Field Description
Route	IP address and network mask for the unicast route
RpfNbr	RPF neighbor
Iif	Incoming interface for the unicast route
Pref	Preference value for the unicast route
Metric	Value of metric for the unicast route (type of metric varies with the unicast protocol)
Count of entries	Number of unicast routes that PIM sparse mode is using

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim sparse-mode unicast-route**

Monitoring the Threshold for Switching to the Shortest Path Tree at a PIM Designated Router for IPv4

Purpose Display the threshold for switching to the shortest path tree at a PIM designated router.

Action To display the threshold for switching to the shortest path tree at a PIM designated router:

```
host1:2#show ip pim spt-threshold
Access List Name                  SptThreshold(in kbps)
-------------------------------------------------------
1                                 infinity
```

Meaning Table 40 on page 136 lists the **show ip pim spt-threshold** command output fields.

Table 40: show ip pim spt-threshold Output Fields

Field Name	Field Description
Access List Name	Name of the IP access list that specifies the groups to which the threshold applies
SptThreshold (in kbps)	Value at which PIM sparse mode switches from a shared tree to an SPT. A value of infinity indicates that PIM sparse mode never switches to an SPT

Related Documentation

- Configuring PIM for IPv4 Multicast on page 89
- **show ip pim spt-threshold**

CHAPTER 7

Configuring DVMRP

E Series routers support Distance Vector Multicast Routing Protocol (DVMRP) on VRs to forward multicast datagrams through a network. DVMRP is an interior gateway protocol that supports operations within an autonomous system, but not between autonomous systems. The multicast backbone of the Internet, MBone, uses DVMRP to forward multicast datagrams. This chapter describes how to configure DVMRP on E Series routers; it contains the following sections:

DVMRP Overview

NOTE: PIM has gained general acceptance among a large number of multicast-enabled networks. We recommend that you use PIM rather than DVMRP for applications that are not otherwise required to run DVMRP.

DVMRP is a dense-mode multicasting protocol and therefore uses a broadcast and prune mechanism. The protocol builds a source-rooted tree (SRT) in a similar way to PIM dense mode. DVMRP routers flood datagrams to all interfaces except the one that provides the shortest unicast route to the source. DVMRP uses pruning to prevent unnecessary sending of multicast messages through the SRT.

A DVMRP router sends prune messages to its neighbors if it discovers that:

- The network to which a host is attached has no active members of the multicast group.
- All neighbors, except the next-hop neighbor connected to the source, have pruned the source and the group.

When a neighbor receives a prune message from a DVMRP router, it removes that neighbor from its (S,G) pair table, which provides information to the multicast forwarding table.

When a host on a previously pruned branch attempts to join a multicast group, it sends an IGMP message to its first-hop router. The first-hop router then sends a graft message upstream.

Identifying Neighbors

In this implementation of DVMRP, a *neighbor* is a directly connected DVMRP router. When you enable DVMRP on an interface, the associated VR adds information about local networks to its DVMRP routing table. The VR then sends probe messages periodically to learn about neighbors on each of its interfaces. To ensure compatibility with other DVMRP routers that do not send probe messages, the VR also updates its DVMRP routing table when it receives route report messages from such routers.

Advertising Routes

As its name suggests, DVMRP uses a distance-vector routing algorithm. Such algorithms require that each router periodically inform its neighbors of its routing table. DVMRP routers advertise routes by sending DVMRP report messages. For each network path, the receiving router picks the neighbor advertising the lowest cost and adds that entry to its routing table for future advertisement.

The cost, or metric, for this routing protocol is the hop count back to the source. The hop count for a network device is the number of routers on the route between the source and that network device.

Table 41 on page 139 shows an example of the routing table for a DVMRP router.

Table 41: Sample Routing Table for a DVMRP Router

Source Subnet	Subnet Mask	From Router	Metric	Time Before Entry Is Deleted from Routing Table	Input Port	Output Port
143.2.0.0	255.255.0.0	143.32.44.12	4	85	3/0	4/0, 4/1
143.3.0.0	255.255.0.0	143.2.55.23	2	80	3/1	4/0, 4/1
143.4.0.0	255.255.0.0	143.78.6.43	3	120	3/1	4/0, 4/1

The DVMRP router maintains an (S,G) pair table that provides information to the multicast forwarding table. The (S,G) pair table is based on:

- Information from the DVMRP routing table
- Information learned from prune messages
- If IGMP and DVMRP are on the same interface, group information learned from IGMP

The (S,G) pair table includes a route from each subnetwork that contains a source to each multicast group of which that source is a member. These routes can be static or learned routes. Table 42 on page 139 shows an example of the (S,G) pair table for DVMRP. A dash (—) in the Input Port column indicates that the interface is associated with a protocol other than DVMRP.

Table 42: Sample DVMRP (S,G) Pair Table

Source Subnet	Multicast Group	Time Before Entry Is Deleted from Routing Table	Input Port	Output Port
143.2.0.0	230.1.2.3	85	3/0	4/0, 4/1
	230.2.3.4	75	3/0	4/0, 4/1
	230.3.4.5	60	3/0	4/1
	230.4.5.6	90	—	4/0
143.3.0.0	230.1.2.3	80	3/1	4/0, 4/1

Related Documentation

- DVMRP Platform Considerations on page 139
- DVMRP References on page 140

DVMRP Platform Considerations

For information about modules that support DVMRP on the ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support DVMRP.

For information about modules that support DVMRP on the E120 and E320 Broadband Services Routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support DVMRP.

Related Documentation

- DVMRP References on page 140

DVMRP References

For more information about DVMRP, see the following resource:

- Distance Vector Multicast Routing Protocol—draft-ietf-idmr-dvmrp-v3-11.txt (April 2004 expiration)

Related Documentation

- DVMRP Platform Considerations on page 139

Enabling DVMRP on a Virtual Router

By default, DVMRP is enabled on the router. To enable DVMRP on a VR:

1. Enable multicast routing.

   host1(config)#ip multicast-routing

2. (Optional) Create a VR or access a VR context.

   host1(config)#virtual-router boston

NOTE: **If you do not specify a VR, you can configure DVMRP on the default router.**

You must enable and configure DVMRP on one or more interfaces for DVMRP to function. See "Activating DVMRP on an Interface" on page 141. You can also set DVMRP limits for the VR; see "Configuring DVMRP Limits" on page 141.

Related Documentation

- DVMRP Overview on page 138
- ip multicast-routing
- virtual-router

Activating DVMRP on an Interface

By default, DVMRP is not activated on an interface. Configuring any DVMRP parameter on an interface automatically activates DVMRP on that interface. You can also activate DVMRP on an interface and use the default parameters.

The **ip dvmrp** command automatically creates and enables DVMRP processing on the current VR. Issuing this command identifies this interface as one that DVMRP owns.

To activate DVMRP on an interface:

- Issue the **ip dvmrp** command in Interface Configuration mode.

  host1:boston(config-if)#**ip dvmrp**

  Use the **no** version to remove DVMRP from an interface.

Related Documentation

- DVMRP Overview on page 138
- ip dvmrp

Configuring DVMRP Limits

You can configure DVMRP and IGMP on the same interface. If you configure DVMRP and IGMP on an interface, the router determines that DVMRP owns the interface.

NOTE: **You cannot configure DVMRP and PIM on the same interface.**

When you have enabled DVMRP processing on a VR, you can configure the following settings for that VR:

- Configuring the Number of Routes that the Virtual Router Records in Warnings on page 141
- Configuring the Maximum Number of DVMRP Routes That the Router Advertises on page 142

Configuring the Number of Routes that the Virtual Router Records in Warnings

The following steps describe how you can set the number of DVRMP routes that the router can record before it generates a system log warning message. The warning alerts you so you can identify routers that are injecting large numbers of routes into the MBone.

To configure the number of routes that the router can record before it generates a system log warning message:

- Issue the **ip dvmrp route-hog-notification** command in Global Configuration mode.

  host1:boston(config)#**ip dvmrp route-hog-notification 5000**

  Use the **no** version to restore the default value, 10,000 routes.

Configuring the Maximum Number of DVMRP Routes That the Router Advertises

To limit the number of routes that the router can advertise on each interface:

- Issue the **ip dvmrp route-limit** command in Global Configuration mode.

  host1:boston(config)#ip dvmrp route-limit 5000

  Use the **no** version to restore the default value, 7000 routes.

Related Documentation

- DVMRP Overview on page 138
- ip dvmrp route-hog-notification
- ip dvmrp route-limit

DVMRP Report Filter Overview

You can configure an interface to accept only reports with routes that appear on a standard IP access list. You can refine the set of accepted routes further, by defining a second access list of neighbors who can supply the specified routes.

For example, suppose you define an access list that specifies that the router accepts only reports for the route 172.16.2.0/24. You then define a second access list that specifies that only neighbors 192.168.1.1 and 193.168.1.1 can supply this route. If neighbor 192.168.2.2 supplies the route, the DVMRP router rejects this report.

You can also modify the value (distance) that the router associates with a DVMRP route when it computes the RPF interface for the source of a multicast packet. By default, the router associates a distance of 0 with DVMRP routes; this value specifies that the router use DVMRP, rather than a unicast routing protocol, to transport multicast datagrams.

However, in a configuration where PIM discovers multicast routes and a unicast routing protocol performs RPF lookups, you can increase the administrative distance to favor the unicast protocol.

For information about defining access lists, see *Configuring Routing Policy* in the *JunosE IP Services Configuration Guide*.

Related Documentation

- DVMRP Overview on page 138
- Filtering DVMRP Reports on page 142

Filtering DVMRP Reports

The following procedure describes how to filter routes in DVMRP reports in accordance with a standard IP access list. You can specify a standard IP access list of sources for which the interface can accept routes, specify a DVMRP administrative distance to favor a unicast routing protocol, or specify a neighbor list to restrict the neighbors from which reports for routes on the first list can be accepted.

To filter routes in a DVMRP report:

- Issue the **ip dvmrp accept-filter** command in Interface Configuration mode.

  host1:boston(config-if)#ip dvmrp accept-filter boston-list 4 neighbor-list boston-neighbors

  Use the **no** version to disable a filter.

Related Documentation

- ip dvmrp accept-filter

DVMRP Summary Addresses Overview

You can configure an interface to advertise a summary address with a known metric rather than a more specific route. DVMRP advertises the summary address if the DVMRP routing table contains a more specific route that matches the address and mask of the summary address.

If you want to advertise all routes rather than a summary, disable automatic summarization on the interface (**no ip dvmrp auto-summary**). By default, the router automatically summarizes DVMRP routes. DVMRP automatic summarization maps a unicast subnet route to a classful network number route when the subnet has a different network number from the IP address of the interface (or tunnel) over which the advertisement travels. If the interface is unnumbered, the router compares the network number of the numbered interface to the IP address to which the unnumbered interface points.

If you configure a summary address on an interface and do not disable automatic summarization, the interface advertises the least-specific address.

Related Documentation

- Configuring DVMRP Summary Addresses on page 143

Configuring DVMRP Summary Addresses

This section describes how you can reenable the router to summarize routes automatically for a specific interface and use it to advertise DVMRP summary addresses on an interface.

To reenable the router to summarize routes automatically for this interface:

- Issue the **ip dvmrp auto-summary** command in Interface Configuration mode..

  host1:boston(config-if)#ip dvmrp auto-summary

  By default, automatic summarization is enabled. Use the **no** version to disable automatic summarization for this interface.

To advertise DVMRP summary addresses on an interface:

- Issue the **ip dvmrp summary-address** command in Interface Configuration mode.

  host1:boston(config-if)#ip dvmrp summary-address 192.48.1.2 255.255.255.0 metric 1

By default, an interface advertises only summary addresses generated by automatic summarization. If you configure multiple overlapping summary addresses on an interface, the one with the shortest mask takes preference. Use the **metric** keyword to specify a DVMRP metric (hop count); the default metric value is 1. Use the **no** version to stop advertising a summary address on the interface.

Related Documentation

- DVMRP Summary Addresses Overview on page 143
- ip dvmrp auto-summary
- ip dvmrp summary-address

Changing the Metric for a Route

The metric for DVMRP is hop count. For example, a route with two hops over a slow serial line is preferable to a route with three hops over a faster optical line.

The router increases the number of DVMRP routes in incoming reports by a default metric of one and in outgoing reports by a default of 0. You can change the metric for an interface to promote or demote the preference for associated routes.

To adjust the number of hops associated with a route:

- Issue the **ip dvmrp metric-offset** command in Interface Configuration mode. This action specifies that the route is more efficient or less efficient than an alternative route.

  ```
  host1:boston(config-if)#ip dvmrp metric-offset in 3
  ```

  Use the **no** version to revert to the default settings: 1 for incoming reports and 0 for outgoing reports.

Related Documentation

- Clearing DVMRP Routes on page 148
- Disabling and Removing DVMRP on page 147
- Enabling DVMRP on a Virtual Router on page 140
- Importing Routes from Other Protocols on page 144
- Specifying Routes to Be Advertised on page 145
- Preventing Dynamic Route Distribution on page 146
- ip dvmrp metric-offset

Importing Routes from Other Protocols

You can import routing information from other protocols into the DVMRP routing table. Only routes that appear in the RPF table can be imported.

NOTE: If you want to use IS-IS, OSPF, or RIP routes, you must make those routes available to multicast protocols. See "Defining Static Routes for Reverse-Path Forwarding" on page 7.

1. Issue the **router dvmrp** command in Global Configuration mode to access Router Configuration mode. You can create and enable DVMRP processing on a VR from the Router Configuration mode.

   ```
   host1:boston(config)#router dvmrp
   ```

   Use the **no** version to remove DVMRP from the VR.

2. Specify a route map using the **route-map** command in Global Configuration mode.

   ```
   host1:boston(config-router)#route-map boston-map atm 3/2
   ```

   Use the **no** version to delete the route map. If you do not specify an interface, it removes the global route map if one exists.

3. Import information from one type of routing domain into another using the **redistribute** command in Router Configuration mode.

   Example—Importing routing information from BGP into DVMRP

   ```
   host1:boston(config-router)#redistribute bgp 100 route-map boston-map
   ```

   Use the **no** version to disable redistribution.

Related Documentation

- Clearing DVMRP Routes on page 148
- Changing the Metric for a Route on page 144
- Disabling and Removing DVMRP on page 147
- Enabling DVMRP on a Virtual Router on page 140
- Specifying Routes to Be Advertised on page 145
- Preventing Dynamic Route Distribution on page 146

- redistribute
- router dvmrp
- route-map

Specifying Routes to Be Advertised

By default, if DVMRP owns an interface, that interface advertises all DVMRP routes it has learned to its neighbors. You can specify the routes that the interface advertises by issuing the **ip dvmrp announce-filter** command in conjunction with a standard IP access list. The IP access list defines the DVMRP routes that are advertised.

To specify the DVMRP routes for an interface to advertise:

- Issue the **ip dvmrp announce-filter** in Interface Configuration mode.

host1:boston(config-if)#**ip dvmrp announce-filter boston-list**

Use the **no** version to enable the interface to advertise all DVMRP routes that it has learned.

Related Documentation

- Clearing DVMRP Routes on page 148
- Changing the Metric for a Route on page 144
- Disabling and Removing DVMRP on page 147
- Enabling DVMRP on a Virtual Router on page 140
- Importing Routes from Other Protocols on page 144
- Preventing Dynamic Route Distribution on page 146

- ip dvmrp announce-filter

Preventing Dynamic Route Distribution

By default, if you make changes to a route map, the router dynamically redistributes the routes in DVMRP. To prevent this dynamic redistribution, use the **disable-dynamic-redistribute** command.

To prevent the dynamic redistribution of routes:

- Issue the **disable-dynamic-redistribute** command in Router Configuration mode.

  host1(config-router)#**disable-dynamic-redistribute**

  There is no **no** version.

Related Documentation

- Clearing DVMRP Routes on page 148
- Changing the Metric for a Route on page 144
- Disabling and Removing DVMRP on page 147
- Enabling DVMRP on a Virtual Router on page 140
- Importing Routes from Other Protocols on page 144
- Specifying Routes to Be Advertised on page 145

- disable-dynamic-redistribute

Exchange of DVMRP Unicast Routes Overview

DVMRP maintains its own unicast routing table, based on distance vector calculations. The routing table defines the best-known distance to each destination and how to get there. The router updates the table by exchanging information with its neighbors. The DVMRP routing table is used solely for RPF lookups.

By default, if DVMRP owns an interface, that interface exchanges DVMRP unicast routes with its neighbors, and you cannot disable the exchange of routes. However, you can enable and disable the exchange of DVMRP unicast routes on interfaces that DVMRP does not own.

When an interface exchanges DVMRP routes, the router obtains routes from DVMRP report messages and stores them in its DVMRP routing table. Other multicast protocols, such as PIM, can then use these routes for RPF lookups. With this feature, PIM can use the DVMRP routing table even when the router is not running DVMRP.

All interfaces, including tunnels, support DVMRP unicast routing. DVMRP tunnels use DVMRP multicast routing to support DVMRP unicast routing.

Related Documentation

- Exchanging DVMRP Unicast Routes on page 147

Exchanging DVMRP Unicast Routes

Use the **ip dvmrp unicast-routing** command to enable the exchange of DVMRP unicast routes on an interface not owned by DVMRP. DVMRP tunnels enable the exchange of IP multicast traffic between routers separated by networks that do not support multicast routing. For information about DVMRP tunnels, see *Configuring IP Tunnels* in the *JunosE IP Services Configuration Guide*.

To enable the exchange of DVMRP unicast routes on an interface not owned by DVMRP:

- Issue the **ip dvmrp unicast-routing** command in Interface Configuration Mode.

  host1:boston(config-if)#**ip dvmrp unicast-routing**

  Use the **no** version to disable the exchange of DVMRP unicast routes on an interface not owned by DVMRP.

Related Documentation

- Exchange of DVMRP Unicast Routes Overview on page 146
- ip dvmrp unicast-routing

Disabling and Removing DVMRP

This topic describes how to disable and remove DVMRP on a virtual router or an interface. You can disable DVMRP on a VR or an interface without removing the configuration. You can also remove DVMRP from a VR or an interface.

Use the **disable** command to disable DVMRP processing on a VR without removing the DVMRP configuration.

To disable DVMRP processing on a VR without removing the DVMRP configuration:

- Issue the **disable** command in Interface Configuration.

  host1:boston(config-router)#**disable**

  Use the **no** version to reenable DVMRP processing on a VR. By default, DVMRP processing is enabled.

Use the **ip dvmrp disable** command to disable DVMRP processing on an interface without removing the DVMRP configuration.

To disable DVMRP processing on an interface without removing the DVMRP configuration:

- Issue the **ip dvmrp disable** command in Interface Configuration mode.

  host1:boston(config-if)#**ip dvmrp disable**

  Use the **no** version to reenable DVMRP processing on an interface.

Related Documentation

- Clearing DVMRP Routes on page 148
- Changing the Metric for a Route on page 144
- Enabling DVMRP on a Virtual Router on page 140
- Importing Routes from Other Protocols on page 144
- Preventing Dynamic Route Distribution on page 146
- Specifying Routes to Be Advertised on page 145
- disable
- ip dvmrp
- ip dvmrp disable
- router dvmrp

Clearing DVMRP Routes

You can clear one or more routes from the DVMRP routing table. However, if you do so, the routes might reappear in the routing table if they are rediscovered. If you do not specify any options, the router removes all routes except those associated with its own interfaces from the DVMRP table. If you specify an IP address but not a subnet mask, the router removes the longest route to that IP address from the DVMRP table. If you specify a subnet mask, the router removes that specific route from the DVMRP table.

To clear DVMRP routes from the routing table:

- Issue the **clear ip dvmrp routes** command in Privileged Exec mode.

  host1:boston#**clear ip dvmrp routes**

  There is no **no** version.

Related Documentation

- Changing the Metric for a Route on page 144
- Enabling DVMRP on a Virtual Router on page 140
- Importing Routes from Other Protocols on page 144
- Preventing Dynamic Route Distribution on page 146
- Specifying Routes to Be Advertised on page 145

- Disabling and Removing DVMRP on page 147
- clear ip dvmrp routes

CHAPTER 8

Monitoring DVMRP

This chapter describes how to monitor DVMRP on E Series routers. It contains the following sections:

Setting a Baseline for DVMRP Statistics

Use the **baseline ip dvmrp** to set the counters for DVMRP statistics to zero, which establishes a reference point, or baseline, for DVMRP statistics.

To set the counters for DVMRP statistics to zero:

- Issue the **baseline ip dvmrp** command in Privileged Exec mode.

  (host1)#**baseline ip dvmrp**

  There is no **no** version.

Related Documentation

- baseline ip dvmrp

Monitoring DVMRP Information for a Virtual Router

Purpose Display DVMRP information for a VR.

Action To display DVMRP information for a VR:

```
host1:boston>show ip dvmrp
Routing Process DVMRP - Distance Vector Multicast Routing Protocol
  Dvmrp Administrative State:         Enabled
  Multicast Administrative State:     Enabled
  Dvmrp Version:                      3.255
  Generation ID:                      0x46828e2b
```

```
Number of Routes:                    2
Number of Triggered Routes:          0
Reachable Routes:                    2
route-hog-notification:              10000
route-limit:                         7000
Send-S32-Prunes-Only:                true
unicastRoutingOnly:                  false
Graceful Restart Duration:           60
Graceful Restart is:                 complete (timer 0 seconds)
Redistribution                       None Configured
dynamic-redistribution:              enabled
```

Meaning Table 43 on page 152 lists the **show ip dvmrp** command output fields.

Table 43: show ip dvmrp Output Fields

Field	Description
Dvmrp Administrative State	State of DVMRP in the software: Enabled or Disabled
Mcast Administrative State	State of multicasting in the software: Enabled or Disabled
Dvmrp Version	Version of DVMRP with which this software is compatible
GenerationID	A number the router generates each time it reboots; when the number changes, neighbors discard all information previously learned from the router
Number of Routes	Number of routes in the DVMRP routing table
Number of Triggered Routes	Number of routes waiting to be advertised, because a parameter for the route changed
Reachable Routes	Number of routes that the router can currently reach
Route-hog-notification	Number of DVMRP routes that the router can record before it generates a system log warning message
Route-limit	Maximum number of routes that the router can advertise on each interface
Send-S32-Prunes-Only	Indicator of whether the router sends only S-32 prunes
true	Router sends only S-32 prunes and grafts to ensure compatibility with other protocols, such as PIM
false	Router sends S-32 and S/Prefix grafts and prunes

Related Documentation

- show ip dvmrp

Monitoring DVMRP Parameters for a Specific Interface

Purpose Display DVMRP parameters for a specified interface.

Action To display DVMRP parameters for a specified interface:

```
host1:v3#show ip dvmrp interface
Interface: ATM2/0.1
      SourceAddress:                      1.0.0.1
      Network/Mask:                    1.0.0.1/24
      Received Bad Packets:                     0
      Received Bad Routes:                      0
      Routes Sent:                              0
      Administrative State:               Enabled
      Summary Address(es)          None Configured
      auto-summary:                       Disabled
      metric-offset in:                         1
      metric-offset out:                        0
      announce-filter:                       None
      accept-filter(s)             None Configured
```

To display a summary of DVMRP parameters for a specified interface:

```
host1:boston#show ip dvmrp interface brief
  Interface          SourceAddress    Network/Mask     RBdPk RBdRt SntRt
  atm5/0.14          14.0.1.1         14.0.1.1/8         0     0     2
  atm5/0.15          15.0.1.1         15.0.1.1/8         0     0     2
```

Meaning Table 44 on page 153 lists the **show ip dvmrp interface** command output fields.

Table 44: show ip dvmrp interface Output Fields

Field	Description
Interface	Type and specifier of the interface connected to a source. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.
SourceAddress	IP address of the interface or, for an unnumbered interface, address of the loopback interface
Network/Mask	Network and mask of the subnet on which the interface resides
Received Bad Packets/RBdPk	Number of bad packets received on this interface
Received Bad Routes/RBdRt	Number of bad routes received on this interface
Routes Sent/SntRt	Number of bad routes advertised on this interface
Administrative State	Configured state of DVMRP on this interface: enabled or Disabled
Summary Address(es)	Specific summary address or addresses that this interface should advertise
auto-summary	Status of automatic summarization: Enabled or Disabled
metric-offset in	Number of hops by which the router increases a DVMRP route advertised in incoming DVMRP reports
metric-offset out	Number of hops by which the router increases a DVMRP route advertised in outgoing DVMRP reports

Table 44: show ip dvmrp interface Output Fields *(continued)*

Field	Description
announce-filter	Routes advertised by the interface
accept-filter(s)	Names of IP access lists that specify the sources for which the interface accepts routes

Related Documentation

- show ip dvmrp interface

Monitoring DVMRP Multicast Groups

Purpose Display information about DVMRP routes to multicast groups.

Action To display information about DVMRP routes to multicast groups:

```
host1:boston#show ip dvmrp mroute
IP DVMRP Multicast Routing Table
(40.0.0.0/16,  228.1.1.1)  Uptime: 77
  Upstream Prune: none
  RPF Interface
      atm5/0.40
  Outgoing interface list:
      atm5/0.31
```

Meaning Table 45 on page 154 lists the **show ip dvmrp mroute** command output fields.

Table 45: show ip dvmrp mroute Output Fields

Field	Description
(S,G) pair	Source, Group pair value
Uptime	Time, in seconds, that this (S, G) pair entry has been in the routing table
Upstream Prune	Whether the router has sent prune messages for this group
RPF Interface	Interface that provides the shortest path back to the source
Outgoing interface list	Types and specifiers of interfaces through which the VR forwards DVMRP messages, such as atm3/0. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.

Related Documentation

- show ip dvmrp mroute

Monitoring DVMRP Neighbors

Purpose Display information about DVMRP neighbors. Specify the brief keyword to view a summary of information.

Action To display information about DVMRP neighbors:

```
host1:boston# show ip dvmrp neighbor
Neighbor Address:                   14.0.0.1
      Interface:                    atm5/0.14
      Neighbor upTime:                28
      Neighbor Major Version:          3
      Neighbor Minor Version:        255
      Neighbor Capabilities:  Prune  GenerationId  Mtrace  NetMask
      Neighbor State:               Active
      Generation ID:        0x3a13fbc2
      Routes Received:                 1
      Bad Routes Received              0
      Bad Packets Received:            0
```

To display a summary of information about DVMRP neighbors:

```
host1:v3#show ip dvmrp neighbor brief
  Interface          NbrAddress      UpTime Maj Min Cap   State
  atm5/0.14          14.0.0.1            32   3 255 PGMN  Active
  atm5/0.15          15.0.0.1            34   3 255 PGMN  Active
```

Meaning Table 46 on page 155 lists the **show ip dvmrp neighbor** command output fields.

Table 46: show ip dvmrp neighbor Output Fields

Field	Description
Neighbor Address/NbrAddress	IP address of the neighbor
Interface	Interface type and specifier, such as atm3/0. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.
Neighbor upTime/UpTime	Length of time, in seconds, that this router has been a neighbor
Neighbor Major Version/Maj	Major number of the DVMRP version on the neighbor
Neighbor Minor Version/Min	Minor number of the DVMRP version on the neighbor
Neighbor Capabilities/Cap	Capability of the neighbor
Prune/P	Ability to send prune messages
Generation ID	Ability to create a generation ID number
Mtrace/M	Ability to trace multicast routes
Netmask/N	Ability to send prunes and grafts with a network mask address
Neighbor State/State	Status of communications with the neighbor
Active	Router is able to communicate with this neighbor
Down	Neighbor is down

Table 46: show ip dvmrp neighbor Output Fields *(continued)*

Field	Description
Ignoring	Router is not accepting messages from this neighbor
Oneway	Router is receiving messages from the neighbor, but the neighbor does not include the router's IP address in the messages. This state can indicate a starting transition, or a problem.
Generation ID	Number that the neighbor generates each time it boots; when the number changes, the VR discards all information previously learned from the router.
Routes Received	Number of routes received from this neighbor
Bad Routes Received	Number of bad routes received from this neighbor
Bad Packets Received	Number of bad packets received from this neighbor

Related Documentation

- show ip dvmrp neighbor

Monitoring DVMRP Routes

Purpose

Display information about DVMRP routes.

You can specify an IP address to display the best route to that address or specify an IP address and subnet mask to display the route that exactly matches this IP address and subnet mask. You can use the **brief** keyword to view a summary of information. You can specify an interface type and specifier to display routes associated with that interface. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.

Action

To display information about DVMRP routes:

```
host1:boston>show ip dvmrp route
Prefix/Length        usNbr/Owner      Metric ExpireTime UpTime Interface
  14.0.0.0/8         Dvmrp Local        1      Never      18    atm5/0.14
    Downstream Interface(s)
        Interface
        atm5/0.15
  15.0.0.0/8         Dvmrp Local        1      Never      18    atm5/0.15
    Downstream Interface(s)
        None
  25.0.0.0/8         14.0.0.1           2        129      11    atm5/0.14
    Downstream Interface(s)
        Interface
        atm5/0.15
```

To display a summary of information about DVMRP routes:

```
host1:v3# show ip dvmrp route brief
  Prefix/Length      usNbr/Owner      Metric ExpireTime UpTime Interface
  14.0.0.0/8         Dvmrp Local        1      Never      26    atm5/0.14
```



```
15.0.0.0/8         Dvmrp Local          1       Never      26    atm5/0.15
25.0.0.0/8         14.0.0.1             2         121      19    atm5/0.14
```

Meaning Table 47 on page 157 lists the **show ip dvmrp route** command output fields.

Table 47: show ip dvmrp route Output Fields

Field	Description
Prefix	IP address of the network
Length	Length of the subnet mask for the network
usNbr/Owner	IP address of the upstream neighbor associated with this route or a description of the origin of the route
Dvmrp Local	Route is associated with a directly attached network
Dvmrp Aggregate	Route is an aggregate route determined by summarization
Metric	Metric associated with this interface for this route
ExpireTime	Time, in seconds, until the VR starts the process for removing the route
UpTime	Length of time, in seconds, that the route has been in the DVMRP routing table
Interface	Type and specifier for the interface, such as atm3/0.

Related Documentation

- show ip dvmrp route

Monitoring the Next Hops of DVMRP Routes

Purpose Display information about the next hop.

Action To display information about the next hop:

```
host1:boston>show ip dvmrp routeNextHop
  addr/mlen          ifIndex  Type
  172.16.0.0/16            4  leaf
  172.17.0.0/16            4  leaf
  172.18.0.0/16            3  leaf
  172.19.0.0/16            3  leaf
  172.19.0.0/16            4  branch
```

Meaning Table 48 on page 157 lists the **show ip dvmrp routeNextHop** command output fields.

Table 48: show ip dvmrp routeNextHop Output Fields

Field	Description
addr	IP address of the next-hop router

Table 48: show ip dvmrp routeNextHop Output Fields *(continued)*

Field	Description
mlen	Mask length of the next-hop router
ifIndex	SNMP interface index for the interface that connects to the next hop
Type	Description of the next-hop router
leaf	Neighbor with no downstream neighbors
branch	Neighbor with downstream neighbors

Related Documentation

- show ip dvmrp routeNextHop

PART 2

Internet Protocol Version 6

CHAPTER 9

Configuring IPv6 Multicast

IPv6 multicast enables a device to send packets to a group of hosts rather than to a list of individual hosts. This chapter describes how to configure IPv6 multicast on the E Series router; it contains the following sections:

IPv6 Multicast Overview

IPv6 defines three types of addresses: *unicast*, *anycast*, and *multicast*. Each type of address enables a device to send datagrams to selected recipients:

- A unicast address enables a device to send a datagram to a single recipient.
- An anycast address enables a device to send a datagram to one recipient out of a set of recipients.

- A multicast address enables a device to send a datagram to a specified set of hosts, known as a multicast group, in different subnetworks.

IPv6 multicast improves network efficiency by allowing a host to transmit a datagram to a targeted group of receivers. For example, a host may want to send a large video clip to a group of selected recipients. It would be time-consuming for the host to unicast the datagram to each recipient individually. If the host broadcasts the video clip throughout the network, network resources are not available for other tasks. The host uses only the resources it needs when multicasting the datagram.

Routers use multicast routing algorithms to determine the best route and transmit multicast datagrams throughout the network. E Series routers support a number of IPv6 multicast protocols on virtual routers. Each virtual router handles the interoperability of IPv6 multicast protocols automatically. To start IPv6 multicast operation on a virtual router, you access the context for that virtual router and configure the desired protocols on the selected interfaces. Table 49 on page 162 describes the function of each the protocol that the router supports.

Table 49: Function of Multicast Protocols on a Router

Protocol	Function
Multicast Listener Discovery (MLD)	Discovers hosts that belong to multicast group.
Protocol Independent Multicast Protocol (PIM)	Discovers other multicast routers that should receive multicast packets.
BGP Multicast Protocol	Routes multicast datagrams between autonomous systems.

The router supports up to 16,384 multicast forwarding entries (multicast routes) at any time.

Reverse-Path Forwarding

IP multicasting uses reverse path forwarding (RPF) to verify that a router receives a multicast packet on the correct incoming interface. The RPF algorithm enables a router to accept a multicast datagram only on the interface from which the router sends a unicast datagram to the source of the multicast datagram.

When the router receives a multicast datagram from a source for a group, the router verifies that the packet was received on the correct RPF interface. If the packet was not received on the correct interface, the router discards the packet. Only packets received on the correct RPF interface are considered for forwarding to downstream receivers.

When operating in sparse-mode, the routers perform an RPF lookup to identify the upstream router from which to request the data and then send join messages for the multicast stream only to that router.

When operating in dense-mode, routers that have multiple paths to the source of the multicast stream initially receive the same stream on more than one interface. In this case, the routers perform an RPF lookup to identify multicast data streams that are not arriving on the best path and send prune messages to terminate these flows.

The RPF lookup need not always be towards the source of the multicast stream. The lookup is done towards the source only when the router is using a source-rooted tree to receive the multicast stream. If the router uses a shared tree instead, the RPF lookup is toward a rendezvous point and not toward the source of the multicast stream.

Multicast Packet Forwarding

Multicast packet forwarding is based on the source (S) of the multicast packet and the destination multicast group address (G). For each (S,G) pair, the router accepts multicast packets on an incoming interface (IIF), which satisfies the RPF check (RPF-IIF). The router drops packets received on IIFs other than the RPF-IIF and notifies the routing protocols that a packet was received on the wrong interface.

The router forwards packets received on the RPF-IIF to a list of outgoing interfaces (OIFs). The list of OIFs is determined by the exchange of routing information and local group membership information. The router maintains mappings of (S,G, IIF) to {OIF1, OIF2...} in the multicast routing table.

You can enable two or more multicast protocols on an IIF. However, only one protocol can forward packets on that IIF. The protocol that forwards packets on an IIF *owns* that IIF. A multicast protocol that owns an IIF also owns the (S,G) entry in the multicast routing table.

Related Documentation

- IPv6 Multicast Platform Considerations on page 163
- IPv6 Multicast References on page 164
- Switch Fabric Bandwidth Configuration on page 6

IPv6 Multicast Platform Considerations

For information about modules that support IPv6 multicasting on the ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support IPv6 multicasting.

For information about modules that support IPv6 multicasting on the E120 and E320 Broadband Services Routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support IPv6 multicasting.

Related Documentation

- IPv6 Multicast Overview on page 161
- IPv6 Multicast References on page 164

IPv6 Multicast References

For more information about IPv6 multicast, see the following resource:

- A " traceroute" Facility for IP Multicast—draft-ietf-idmr-traceroute-ipm-07.txt (January 2001 expiration)

Related Documentation

- IPv6 Multicast Overview on page 161
- IPv6 Multicast Platform Considerations on page 163

Before You Begin

Before you begin configuring multicast on IPv4 and IPv6 interfaces, you must:

- Configure IP interfaces. For more information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*.
- Configure IPv6 interfaces. For more information about configuring IPv6 interfaces, see *Configuring IPv6* in *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- Configuring IPv6 Multicast Attributes on page 165
- Example: Configuring an IPv6 Multicast Bandwidth Map on page 171
- Activating IPv6 Multicast QoS Adjustment Functions on page 175
- Configuring IPv6 Hardware Multicast Packet Replication on page 181
- Blocking IPv6 Mroutes on page 182
- Enabling Interface-Level Admission Bandwidth Limitation for IPv6 on page 183
- Creating IPv6 Mroute Port Limits on page 185
- Enabling Port-Level Admission Bandwidth Limitation for IPv6 on page 185
- Deleting IPv6 Multicast Forwarding Entries on page 187

Switch Fabric Bandwidth Configuration

By default, the switch fabric for the ERX1440, ERX310, E120, and E320 Broadband Services Routers uses a bandwidth weighting ratio of 15:2 for multicast-to-unicast weighted round robin (WRR). In the absence of strict-priority traffic, and when both unicast and multicast traffic compete for switch fabric bandwidth, the switch fabric allocates 15/17ths of the available bandwidth to multicast traffic and 2/17ths of the available bandwidth to unicast traffic.

You can use the **fabric weights** command to change the ratio for multicast to unicast traffic on the router switch fabric. For more information about the **fabric weights** command, see *Configuring the Switch Fabric Bandwidth* in the *JunosE System Basics Configuration Guide*.

Related Documentation

- IPv4 Multicast Overview on page 4
- IPv6 Multicast Overview on page 161

- fabric weights

Configuring IPv6 Multicast Attributes

You can configure IPv6 multicast to control the flow of multicast traffic with the following tasks:

- Enabling IPv6 Multicast on page 165
- Defining IPv6 Static Routes for Reverse-Path Forwarding on page 165
- Enabling and Disabling RPF Checks for IPv6 on page 166
- Specifying Unicast Routes for RPF in IPv6 on page 166
- Defining Permanent IPv6 Multicast Forwarding Entries on page 166

Enabling IPv6 Multicast

IPv6 multicast works on virtual routers. By default, IPv6 multicast is disabled on a virtual router. To enable IPv6 multicast on a virtual router, access the context for a virtual router, and then issue the **ipv6 multicast-routing** command.

To enable IPv6 multicast routing on the default virtual router:

- Issue the **ipv6 multicast-routing** command in Global Configuration mode.

```
host1(config)#ipv6 multicast-routing
```

Use the **no** version to disable IPv6 multicast routing on the virtual router (the default). In the disabled state, all multicast protocols are disabled, and the virtual router forwards no multicast packets.

Defining IPv6 Static Routes for Reverse-Path Forwarding

You can use the **ipv6 rpf-route** command to define RPF to verify that a router receives a multicast packet on the correct incoming interface.

To customize the static IPv6 routes that the router may use for RPF:

- Issue the **ipv6 rpf-route** command in Global Configuration mode.

```
host1(config)#ipv6 rpf-route 1000::/64 ATM2/1.200
```

Use the **no** version to remove the static route.

Enabling and Disabling RPF Checks for IPv6

By default, the router accepts multicast packets for each (S,G) pair on an IIF, which satisfies the RPF check (RPF-IIF). When the router performs RPF checks, only the interface that first accepts traffic for an (S,G) pair accepts subsequent traffic for that pair. If traffic stops coming on that interface and starts arriving on another interface, the router does not accept or forward the traffic.

Some network configurations require the router to accept traffic on any interface. To do so, you can disable the RPF check on a specified set of (S,G) pairs by issuing the **ipv6 multicast-routing disable-rpf-check** command.

When you disable RPF checks, the router accepts multicast packets for (S,G) pairs on any incoming interface. When the router has added the new route to its multicast routing table, it accepts multicast packets for these pairs on any interface in the virtual router and forwards them accordingly. Multicast routes established before you issue this command are not affected.

To enable and disable RPF checks:

- Enable RPF checks for all (S,G) pairs (the default situation).

  host1(config)#**no ipv6 multicast-routing disable-rpf-check**

- Disable RPF checks for the (S,G) pair specified using a standard IPv6 access list.

  host1(config)#**ipv6 multicast-routing disable-rpf-check denver-list**

Specifying Unicast Routes for RPF in IPv6

You can use the **ipv6 route-type** command to specify that BGP or OSPF IPv6 routes should be available for RPF. Routes available for RPF appear in the multicast view of the routing table.

To specify that BGP IPv6 routes are available for both unicast protocols and multicast protocols to perform RPF checks:

host1(config)#**router bgp**
host1(config-router)#**ipv6 route-type both**

Defining Permanent IPv6 Multicast Forwarding Entries

An mroute is a multicast traffic flow, a (S,G) entry used for forwarding multicast traffic. By default, forwarding mroutes (with a valid RPF incoming interface) are timed out if data for them is not received for 210 seconds. However, you can specify an mroute as permanent by using the **ipv6 multicast-routing permanent-mroute** command.

To specify that any newly created mroutes that match the specified access-list do not time out:

- Issue the **ipv6 multicast-routing permanent-mroute** command in Global Configuration mode.

  host1(config)#**ipv6 multicast-routing permanent-mroute routesv61**

Use the **no** version to prevent any new mroutes from becoming permanent. To remove existing permanent mroutes, use the **clear ipv6 mroute** command.

NOTE:

- The **ipv6 multicast-routing permanent-mroute** command does not change existing mroutes.
- Permanent mroutes are removed if a topology change occurs that affects the mroute.
- Permanent mroutes may be removed due to certain protocol actions (for example, PIM sparse mode switching from shared to shortest path tree).
- Outgoing interface lists of permanent mroutes may change due to protocol actions.

Related Documentation

- IPv6 Multicast Overview on page 161
- Before You Begin on page 164
- Deleting IPv6 Multicast Forwarding Entries on page 187

- clear ipv6 mroute
- ipv6 multicast-routing
- ipv6 multicast-routing disable-rpf-check
- ipv6 multicast-routing permanent-mroute
- ipv6 route-type
- ipv6 rpf-route

IPv6 Multicast Bandwidth Map Overview

Multicast interface-level admission control, port-level admission control, and QoS adjustment all use a single multicast bandwidth map. The multicast bandwidth map is a route map that uses the **set admission-bandwidth**, **set qos-bandwidth**, **set admission-bandwidth adaptive**, or **set qos-bandwidth adaptive** commands. The **adaptive** commands configure an auto-sense mechanism for measuring the multicast bandwidth.

NOTE: Even though you can include any of the above commands several times in a route map entry, only the last admission-bandwidth command or qos-bandwidth command in the bandwidth map is used. In other words, if you included the **set qos-bandwidth** command first and then the **set qos-bandwidth adaptive** command, the bandwidth map would use the **set qos-bandwidth adaptive** command.

Interface- and port-level admission control is performed when an outgoing interface (OIF) on the interface or port is added to the mroute for a given (S,G) multicast data stream and the multicast bandwidth map contains a **set admission-bandwidth** or **set admission-bandwidth adaptive** action for that (S,G).

QoS adjustment is performed on the joining interface when an OIF is added to the mroute for a given (S,G) data stream and the multicast bandwidth map contains a **set qos-bandwidth** or **set qos-bandwidth adaptive** action for that (S,G).

NOTE: You can create a single route map with the **set admission-bandwidth** command, the **set qos-bandwidth** command, or both. However, creating an entry with only one of these **set** commands enables only that specific function for the matched address (that is, only multicast traffic admission control or only QoS adjustment). The same is true for the **adaptive** commands.

Related Documentation

- IPv6 Multicast Overview on page 161
- Autosense Mechanism Overview on page 9
- Adaptive Mode Mechanism Overview on page 168
- Example: Configuring an IPv6 Multicast Bandwidth Map on page 171

- set admission-bandwidth
- set qos-bandwidth

Autosense Mechanism Overview

Video bandwidth is typically considered to be a constant rate—2 Mbps for standard definition television (SDTV) and 10 Mbps for high definition television (HDTV). However, in reality, and depending on achievable video compression, the bit rate can vary. For example, HDTV streams (using MPEG4 or WM9 encoding) can vary between 6 Mbps (for low-action programs) to 10 Mbps (for a fast-paced, high-action programs). The autosense mechanism causes the bandwidth value, used for admission control and QoS adjustment, to be the actual measured rate of the stream. Using this feature to measure the actual bandwidth avoids the need to configure arbitrary bandwidth limits and enables a channel to be reassigned to a different (S, G) without requiring a bandwidth map to be changed.

Related Documentation

- Adaptive Mode Mechanism Overview on page 12
- Defining a Multicast Bandwidth Map on page 8

Adaptive Mode Mechanism Overview

You can configure the auto-sense mechanism in the multicast bandwidth using the **set admission-bandwidth adaptive** command, **set qos-bandwidth adaptive** command, or both. For example:

```
host1(config)#route-map mcast-bandwidths permit 10
host1(config-route-map)#match ipv6 address sdtv
host1(config-route-map)#set admission-bandwidth adaptive
host1(config-route-map)#set qos-bandwidth adaptive
host1(config-route-map)#end
```

In this example, any stream with an (S,G) that matches the sdtv access list performs adaptive bandwidth detection for admission control and QoS adjustment.

A rate measurement mechanism runs on the ingress line card that polls the forwarding controller (FC) to obtain statistics for each mroute. This mechanism then reports the rate measurement to the switch route processor (SRP) to update the bandwidth map. By computing the average bandwidth over a relatively short sampling period (T1; 5 seconds), the measurement approximates the peak bandwidth of the multicast stream.

As an example, assume that a new mroute (S1, G1) is added to the interface controller (IC) at time t0.

Figure 14: Example of Adaptive IPv6 Multicast Bandwidth Detection

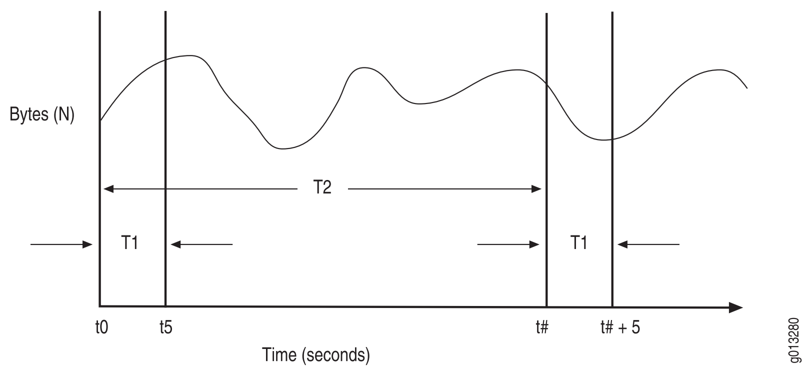

To calculate the measured bandwidth of a stream, the router uses the following equation:

$R = (N_{t+5} - N_{t}) / 5$

Where

R = Calculated bandwidth of the stream during each sampling interval

N_t = Bytes measured at the start of each sampling period (t seconds)

N_{t+5} = Bytes measured at the end of each sampling period (t+5 seconds)

NOTE: When the mroute is first installed in the FC (at t = 0), R_0 is undetermined. For multicast admission control no joins are admitted until the first bandwidth measurement is computed (that is, for admission control, R_0 is considered to be infinite). Similarly, no QoS adjustment occurs until the first bandwidth measurement is computed (that is, for QoS adjustment, R_0 is considered to be zero [0]).

Using the earlier graph as a reference, the first bandwidth rate ($R1_0$) and at time t_5 (N_5) and the bytes received values are subtracted and divided by the time period T_1 to yield the average rate. This process is repeated every sampling interval, T_2, to yield rates R1, R2, R3, and so on.

The first two sampling interval calculations would look like the following:

$R_1 = (N_5 - N_0)/5$

$R_2 = (N_{\#+5} - N_{\#})/5$

The router maintains a history of bandwidth measurements (H) for each mroute, up to a maximum of M measurements. The actual rate, R, reported to the SRP is the maximum rate measured in those H samples.

In order to minimize the IC to SRP traffic generated by the rate measurements, the IC reports a bandwidth change only when a newly computed rate (R#) differs from the current rate by a specified threshold. When R_s is computed at time t = 5 seconds, R is set to R_1. A rate update occurs whenever a newly calculated rate (R) differs from R_1 by at least a threshold value (specified as a percentage, P) of the measured peak bandwidth. This calculation would look like the following:

R = Rt, if and only if the absolute value of (R - Rt) > P * R.

The values assigned to variables associated with this algorithm are as shown in Table 50 on page 170.

Table 50: Adaptive Mode Algorithm Values

Variable	Value	Units	Description
T1	5	Seconds	Sampling period; the time in which a sample is taken
T2	0	Seconds	Sampling interval; zero (0) seconds indicates continuous sampling
H	12	Samples	Number of history samples over which to compute measurement
M	12	Samples	Maximum number of samples maintained in history
P	1	Percent	Threshold value; percent difference by which a newly calculated rate must differ from the measured peak bandwidth before a rate update occurs

Related Documentation

- IPv6 Multicast Bandwidth Map Overview on page 167
- Autosense Mechanism Overview on page 9
- Example: Configuring an IPv6 Multicast Bandwidth Map on page 171

- match ipv6 address
- route-map

- set admission-bandwidth
- set qos-bandwidth

Example: Configuring an IPv6 Multicast Bandwidth Map

This example shows how to create a multicast bandwidth map for both multicast traffic admission control and QoS adjustment.

- Requirements on page 171
- Overview on page 171
- Configuring an IPv6 Multicast Bandwidth Map on page 172

Requirements

This example uses the following hardware and software components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Before you begin configuring multicast on IPv6 interfaces, you must:

- Configure IPv6 interfaces. For more information about configuring IPv6 interfaces, see *Configuring IPv6* in *JunosE IP, IPv6, and IGP Configuration Guide*.

Overview

The multicast bandwidth map is a route map that uses the **set admission-bandwidth**, **set qos-bandwidth**, **set admission-bandwidth adaptive**, or **set qos-bandwidth adaptive** commands. Multicast interface-level admission control, port-level admission control, and QoS adjustment all use a single multicast bandwidth map.

The **adaptive** commands configure an auto-sense mechanism for measuring the multicast bandwidth.

Configuring an IPv6 Multicast Bandwidth Map

Configuring a Route Map

Step-by-Step Procedure

Define a route map using the **set admission-bandwidth** and **set qos-bandwidth** commands.

NOTE: In this example, you can replace the **set admission-bandwidth** command and **set qos-bandwidth** command with their **adaptive** command counterparts.

1. Define a route map.

   [edit]
   host1(config)#**route-map mcast-bandwidths permit 10**

2. Match the route map to an access list.

   [edit]
   host1(config-route-map)#**match ipv6 address sdtv**

3. Configure multicast bandwidths.

   [edit multicast bandwidths for admission control and QoS adjustment]
   host1(config-route-map)#**set admission-bandwidth 2000000**
   host1(config-route-map)#**set qos-bandwidth 2000000**

4. Configure the route map.

   [edit]
   host1(config-route-map)#**route-map mcast-bandwidths permit 20**

5. Match the route map to an access list.

   [edit]
   host1(config-route-map)#**match ipv6 address hdtv**

6. Configure multicast bandwidths.

   [edit multicast bandwidths for admission control and QoS adjustment]
   host1(config-route-map)#**set admission-bandwidth 10000000**
   host1(config-route-map)#**set qos-bandwidth 10000000**

Configuring an Access List

Step-by-Step Procedure

Define the access list for use by the **match ipv6 address** command to match (S,G) and (*,G) entries.

NOTE: You can also define a prefix-list or a prefix-tree for use by the **match ipv6 address** command to match (S,G) and (*,G) entries.

1. Configure access lists.


```
[edit access lists]
host1(config)#access-list sdtv permit ip host 31::1 ff3e::0/112
host1(config)#access-list hdtv permit ip host 32::1 ff3e::0/112
host1(config)#access-list hdtv permit ip host 32::2 ff3e::0/112
```

Related Documentation

- IPv6 Multicast Bandwidth Map Overview on page 167
- *Configuring Routing Policy* in *JunosE IP Services Configuration Guide*
- access-list
- match ipv6 address
- route-map
- set admission-bandwidth
- set qos-bandwidth

Multicast QoS Adjustment for IPv6

When the router uses multicast outgoing interface (OIF) mapping, any multicast streams that a subscriber receives bypass any configured QoS treatment for that subscriber interface. The Multicast QoS adjust feature provides a way in which the router can account for this multicast traffic.

Multicast OIF Mapping Case for IPv6

Multicast outgoing interface (OIF) mapping enables the router to decrease the inefficiencies associated with replicating streams of multicast traffic. Using OIF maps, Multicast Listener Discovery (MLD) joins that the router receives on a subscriber interface can be mapped to a special interface for forwarding. This special interface can be on a different physical port or line module from that of the join interface.

NOTE: For additional information about how to configure OIF mapping, see “Configuring Group Outgoing Interface Mapping” on page 62.

Using this mapping function, the router can send a single copy of each multicast stream over the special interface and the access nodes are configured to perform any final replication to the subscribers and merge unicast and multicast data flows onto the subscriber interfaces as necessary. See Figure 15 on page 174.

Figure 15: Multicast OIF Mapping

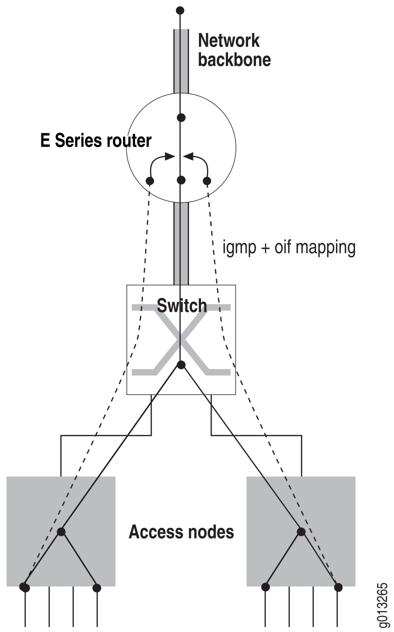

One disadvantage to using multicast OIF mapping is that the multicast traffic bypasses any QoS treatment that is applied to subscriber interfaces. Configuring QoS adjustment resolves this problem. (See Parameter Definition Attributes for QoS Administrators Overview for additional information about configuring QoS adjustment.) With QoS adjustment configured, when a subscriber requests to receive a multicast stream (or, more appropriately, when an OIF is added to the mroute), the router reduces the unicast QoS bandwidth applied to the subscriber interface (that is, the join interface) by the amount of bandwidth for that multicast stream.

Multicast Traffic Receipt Without Forwarding for IPv6

In this case, the router is not given the responsibility of forwarding multicast streams. Instead, the service provider arranges for the router to receive the multicast streams so the router can detect the flow and perform QoS adjustment. An outgoing interface (OIF) map is installed that maps the traffic streams to a loopback interface configured for Multicast Listener Discovery (MLD) version passive. This means that when the traffic is received, a null mroute is installed (that is, an mroute with an empty OIF list) and the router applies the QoS adjustment to the join interface. See Figure 16 on page 175.

NOTE: Ensure that Protocol Independent Multicast-single mode (PIM-SM) (or any other upstream multicast protocol) is informed of the group (or source-group) interest.

Figure 16: Multicast Traffic Receipt Without Forwarding

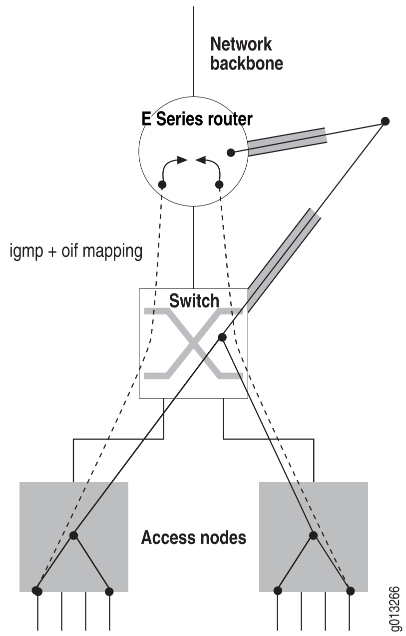

Activating IPv6 Multicast QoS Adjustment Functions

The **ipv6 multicast-routing bandwidth-map** command activates the specified bandwidth map. By activating the bandwidth map, this command also activates the multicast QoS adjustment function contained in the bandwidth map.

CAUTION: To activate multicast QoS adjustment, you must first create a bandwidth map. See "Example: Configuring an IPv6 Multicast Bandwidth Map" on page 171 for details.

To enable the QoS adjust function on the router with the configured route map:

- Issue the **ipv6 multicast-routing bandwidth-map** command in Global Configuration mode.

  host1(config)#**ipv6 multicast-routing bandwidth-map mcast-bandwidths**

Use the **no** version to disable the multicast QoS adjustment function on the router.

Related Documentation

- Before You Begin on page 164
- Multicast OIF Mapping Case for IPv6 on page 173
- Multicast Traffic Receipt Without Forwarding for IPv6 on page 174
- Activating IPv6 Multicast QoS Adjustment Functions on page 175

- ipv6 multicast-routing bandwidth-map

IPv6 Hardware Multicast Packet Replication Overview

You can configure IPv6 multicast to replicate packets to optimized hardware on a logical port instead of using the forwarding controller (FC) on the router.

The bandwidth between the line module and the I/O module or input/output adapter (IOA) on the E Series router is limited. A high-density Ethernet module provides eight physical ports that can consume the bandwidth between the line module and the I/O module or IOA before providing enough traffic to support egress line rate for all of these ports.

Figure 17 on page 176 displays how multicast traffic is typically replicated on the line module. Each of these replicated packets is transmitted from the line module to the I/O module or IOA.

Figure 17: Packet Flow Without Hardware Multicast Packet Replication

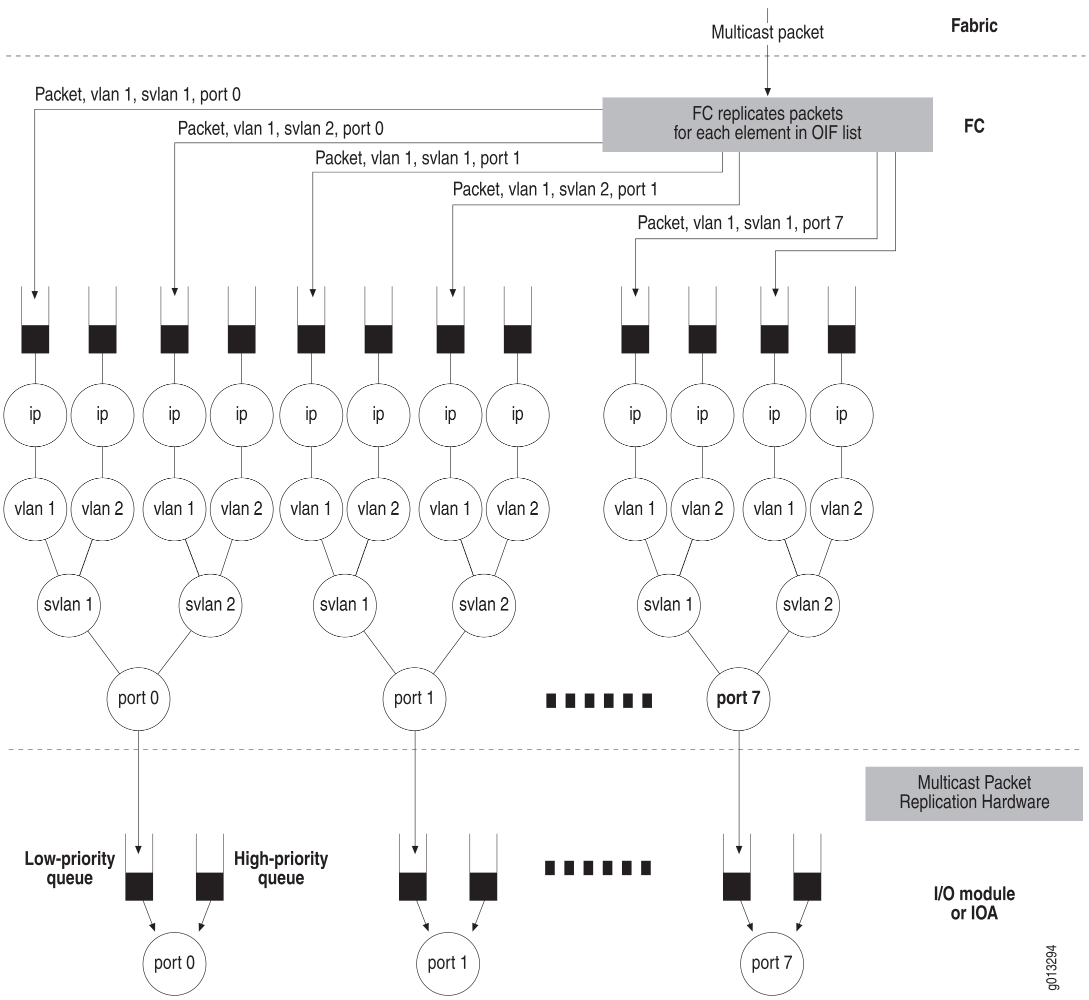

The hardware multicast packet replication feature enables you to configure multicast traffic for a virtual LAN (VLAN) or service VLAN (S-VLAN) to be replicated on the I/O module or IOA so that only one copy of the packet is transmitted from the line module to the I/O module or IOA. Replication for each of the ports is performed on the I/O module or IOA.

Configuring hardware multicast packet replication for high-density Ethernet is useful when you want to provide the same multicast stream out of some or all of the ports, such as for IP television (IPTV). Configuring hardware multicast packet replication enables you to:

- Reduce the number of packets sent from the FC to the module.
- Reduce the CPU consumed by the FC processing each elaboration of the packet.

You can use the feature to increase the bandwidth of multicast traffic out of each of the Gigabit Ethernet ports.

Figure 18 on page 177 displays the flow of a multicast packet using the hardware multicast packet feature.

Figure 18: Packet Flow with Optimized Multicast Packet Replication

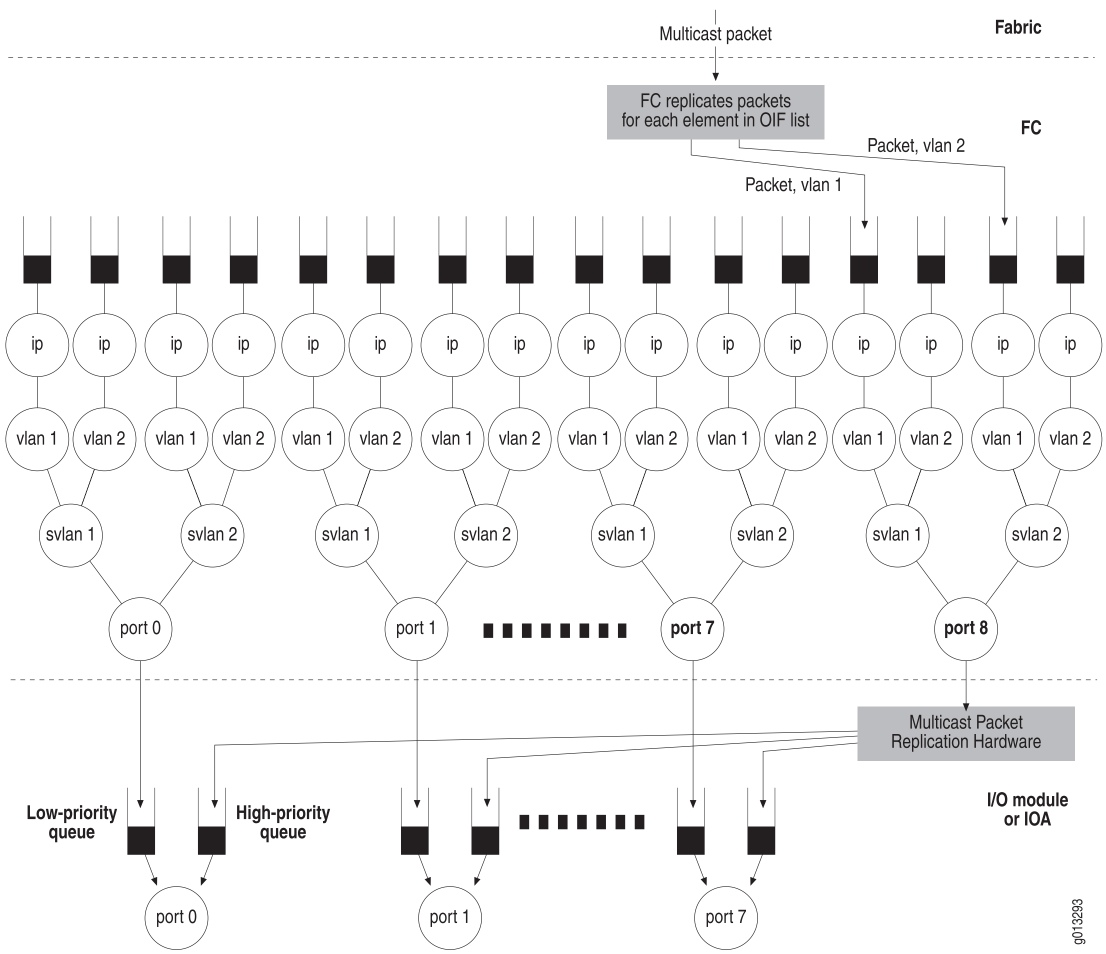

Each high-density Ethernet module has eight physical ports, numbered 0–7. A logical port is available for the hardware multicast packet replication feature, numbered port 8.

JunosE tracks the outgoing interfaces (OIFs) in an mroute that have been redirected to use the hardware multicast packet replication hardware. The system accepts only egress multicast traffic to traverse the interface stack on the enabled port. The system drops unicast traffic that is routed to this port.

Each port on the I/O module or IOA displayed in Figure 18 on page 177 has two queues. These queues are further down the egress path than the queues found on the line module and populated by the FC.

The low-priority queue is dedicated to packets that are received from the line module queues that are dedicated to the physical ports. This queue blocks when full and provides backpressure to the line module. This queue services unicast and multicast traffic that is not using the hardware multicast packet replication feature.

The high-priority queue is dedicated to packets that are received from the line module queue for port 8. This queue is serviced at a higher priority than the first queue, and drops packets when full.

For more information about high-density Ethernet, see *Configuring Ethernet Interfaces* in *JunosE Physical Layer Configuration Guide*.

Supported Modules and Encapsulations

You can enable optimized multicast packet replication on port 8 of the following high-density Ethernet modules:

- GE-8 I/O module (pairs with the GE-HDE line module)
- ES2-S1 GE-8 IOA (pairs with the ES2 4G LM and the ES2 10G LM)

When enabled, the optimized multicast packet replication feature defines the encapsulation of the egress multicast packet. The following encapsulations are supported:

- IPv6 over Gigabit Ethernet
- IPv6 over VLAN
- IPv6 over S-VLAN

NOTE: **802.3ad link aggregation group (LAG) bundles do not support optimized multicast packet replication.**

The optimized multicast packet replication feature also provides an interface over which you can configure the following:

- IP maximum transmission unit (MTU)
- Ethernet MTU

- Egress IP policy
- Egress VLAN policy
- QoS

Relationship with OIF Mapping

Multicast OIF mapping enables the router to decrease the inefficiencies associated with replicating streams of multicast traffic. Using OIF maps, Multicast Listener Discovery (MLD) joins that the router receives on a subscriber interface can be mapped to a special interface for forwarding.

The hardware multicast packet replication feature enables you to redirect each of the IPv6 interfaces on a line module over a dedicated multicast VLAN to a single IPv6 interface over port 8. The FC is only required to send a single packet per dedicated multicast VLAN to the I/O module or IOA. The module then replicates this packet to the appropriate ports.

For more information about configuring OIF mapping, see "Configuring MLD Group Outgoing Interface Mapping" on page 242.

Related Documentation

- IPv6 Hardware Multicast Packet Replication Considerations on page 179
- Configuring IPv6 Hardware Multicast Packet Replication on page 181

IPv6 Hardware Multicast Packet Replication Considerations

When configuring hardware multicast packet replication, the following considerations apply.

- Do not configure or transmit routing protocols over port 8. The FC drops traffic routed to an IPv6 interface stacked over port 8.
- We recommend that you configure the IP address of the IPv6 interface over port 8 to be unnumbered.
- We recommend that you configure an IPv6 interface over a VLAN over one of the physical ports to reference the IPv6 interface over the same VLAN over port 8.

  You cannot create the following configurations:

  - When two IPv6 interfaces configured over a port reference the same IPv6 interface over port 8. The system does not accept this configuration attempt because you typically configure the hardware multicast packet replication feature to redirect multicast traffic over one VLAN, then redirect it to the same VLAN on port 8.
  - When the IPv6 interface configured with the hardware multicast packet replication attribute is not installed on a line module that supports hardware multicast packet replication.

  - When the IPv6 interface designated by the hardware multicast packet replication attribute is not installed on a line module that supports hardware multicast packet replication.
  - When the IPv6 interface designated by the hardware multicast packet replication attribute is not on the same line module as the IPv6 interface configured with this attribute.

- When you configure a unique source MAC address for VLANs on port 8, the hardware multicast packet replication hardware stamps the source MAC address on the VLAN, overwriting any MAC address that you configured. For more information, see *Configuring Ethernet Interfaces* in *JunosE Physical Layer Configuration Guide*.
- The regular multicast implementation utilizes interface stacking that provides a unique IPv6 attachment point for each elaboration of the egress multicast packet.

  For the hardware multicast packet replication feature, you must attach policies to an interface stack over port 8 that defines the encapsulation of the egress multicast traffic. The system supports policies over port 8 just as it is above any of the other ports on this line module.

  Policies applied to the interface stack over port 8 affect the packets traversing this stack whether or not the packet is destined for one port or all of the physical ports. Therefore, you cannot apply different egress policies to multicast traffic for the interfaces stacked above different ports, or rate limit on an individual interface over a port. You also cannot monitor policy statistics on individual interfaces over a port.

  Instead, you can apply egress policy to an interface stacked over port 8. The system applies the policy before the packet has been elaborated for each of the ports.

- The JunosE QoS component provides hierarchical egress scheduling and shaping on Gigabit Ethernet ports 0–7. The regular multicast implementation replicates packets on the FC, with each replicated packet placed on a line module queue destined for a single physical port. The line module queue can also receive QoS behavior specific to that queue.

  For the hardware multicast packet replication feature, the FC does not replicate the packet for each of the individual ports. Instead, it places the packet on a special queue destined for port 8.

  You can configure QoS on the packets flowing through port 8, but this has limited value because each packet passed through this port can be transmitted through one of more of the physical ports. Therefore, the packets placed on this special queue might not receive the same QoS behavior as ports 0–7.

  We recommend that you configure the network so the I/O or IOA queues are not oversubscribed. The traffic transmitted by the physical port is a combination of packets from the two I/O or IOA queues. When the sum of the packets in these queues is greater than line rate, the system can drop traffic that is not using hardware multicast packet replication.

  When you configure a traffic shaper on a physical port and configure hardware multicast packet replication, the packets created using the feature avoid the traffic shaper for that port. To control this, you can use traffic shaper on the physical port and port 8. The sum of the traffic shapers must be less than or equal to the line rate of the port.

A traffic shaper on port 8 can result in the overall utilization of egress bandwidth for any one port being less the line rate because the packets being replicated might not be transmitted to every port. Packets destined to some of the ports contribute to the traffic shaping for all of the ports on the I/O module or IOA.

Related Documentation

- IPv6 Hardware Multicast Packet Replication Overview on page 176
- Configuring IPv6 Hardware Multicast Packet Replication on page 181

Configuring IPv6 Hardware Multicast Packet Replication

To configure hardware multicast packet replication:

1. Configure port 8 on a high-density Ethernet module to accept redirected egress multicast traffic.
   a. Specify the Gigabit Ethernet interface on port 8.
   b. Create a VLAN major interface.
   c. Create a VLAN subinterface.
   d. Assign a VLAN ID.
   e. Configure an unnumbered IPv6 interface.
   f. Enable MLD on the interface with only multicast-data-forwarding capability.

```
host1(config)#interface gigabitEthernet 2/8
host1(config-if)#encapsulation vlan
host1(config-if)#interface gigabitEthernet 2/8.1
host1(config-if)#vlan id 1
host1(config-if)#ipv6 unnumbered loopback 0
host1(config-if)#ipv6 mld version passive
```

2. Configure an IPv6 interface to redirect egress multicast traffic to port 8.
   a. Create a VLAN subinterface.
   b. Assign a VLAN ID.
   c. Assign an IPv6 address.
   d. Configure the interface to redirect egress multicast traffic to port 8.

```
host1(config)#interface gigabitEthernet 2/0.101
host1(config-if)#vlan id 1
host1(config-if)#ipv6 address 1::1/64
host1(config-if)#ipv6 multicast ioa-packet-replication gigabitEthernet 2/8.1
```

Related Documentation

- Before You Begin on page 164

- encapsulation vlan
- interface gigabitEthernet
- ipv6 address
- ipv6 mld version
- ipv6 multicast ioa-packet-replication
- ipv6 unnumbered
- vlan id

Interface-Level Multicast Traffic Configuration for IPv6

You can block mroute creation and limit the multicast bandwidth admitted on an outgoing interface.

Blocking IPv6 Mroutes

By default, when an interface receives multicast traffic, even when the scope of that traffic exceeds link-local, the virtual router creates an mroute. You can use the **ipv6 block-multicast-sources** command to block all multicast traffic with a scope larger than link-local (for example, global) and prevent mroute creation under these conditions.

NOTE: Issuing this command does not affect reception of link-local multicast packets.

To block all multicast traffic with a scope larger than link-local:

- Issue the **ipv6 block-multicast-sources** command in Global Configuration mode.

  host1(config)#**ipv6 block-multicast-sources**

Use the **no** version to restore the default behavior of creating mroutes on receiving multicast packets.

Interface-Level Admission Bandwidth Limits for IPv6 Overview

Interface-level multicast admission control is performed when an outgoing interface (OIF) on the interface is added to the mroute for a given (S,G) multicast data stream and the multicast bandwidth map contains a **set admission-bandwidth** action for that (S,G). When enabled, the admission-bandwidth for a particular (S,G) is read from the multicast bandwidth map and recorded in the mroute when the (S,G) mroute is created.

CAUTION: **Before you can limit interface-level admission bandwidth, you must first create a bandwidth map. See** "Example: Configuring an IPv6 Multicast Bandwidth Map" on page 171 **for details.**

When an OIF is subsequently added to the mroute, the OIF is blocked from forwarding data if the additional bandwidth contributed by the (S,G) exceeds the admission-bandwidth limit for the interface. In JunosE releases earlier than Release 12.0.x, in an OIF mapping scenario where the digital subscriber line access multiplexer (DSLAM) does not perform per-subscriber multicast admission control, the router disregards the multicast admission bandwidth limit configured on the join interface. If the limit configured on the mapped interface exceeds the admission-bandwidth limit for the interface, the router blocks the mapped interface from forwarding data.

Now, in an OIF mapping scenario where the DSLAM performs per-subscriber multicast admission control, the router checks the bandwidth limit configured on the join interface. If the multicast stream is forwarded over the mapped interface, the router admits the multicast stream and forwards the stream to the join interfaces whose bandwidth does not exceed the configured bandwidth limit. The router also performs QoS Adjust for the multicast stream on the unblocked (forwarding) subscriber interfaces. The router does not replicate the stream to the subscriber interfaces whose bandwidth exceeds the configured bandwidth limit and it does not perform QoS Adjust for the multicast stream on the blocked subscriber interfaces.

If the multicast stream is not forwarded over the mapped interface, the router blocks the multicast stream and does not forward the stream. The router also does not perform QoS Adjust for the multicast stream on the blocked subscriber interface.

Enabling Interface-Level Admission Bandwidth Limitation for IPv6

You can use the **ipv6 multicast admission-bandwidth-limit** command to enable multicast admission control on interfaces (including dynamic IP interfaces) that are configured to run MLD. You can also use this command on a PIM (sparse-mode, dense-mode, or sparse-dense-mode) interface if MLD is configured on the interface (including the **ipv6 mld version passive** command).

To limit bandwidth for an interface that accepts MLD groups:

- Issue the **ipv6 multicast admission-bandwidth-limit** command in Interface Configuration mode.

  ```
  host1:boston(config-if)#ipv6 multicast admission-bandwidth-limit 2000000
  ```

Use the **no** version to remove the bandwidth limitation for the interface.

OIF Interface Reevaluation for IPv6

If you change the admission bandwidth for an interface, all mroutes with that interface as an outgoing interface (OIF) are reevaluated as follows:

- If the bandwidth limit is increased, blocked OIFs may become unblocked. If the interface is a blocked OIF on multiple mroutes, the order in which the mroutes are visited, and which (S,G) streams become unblocked, is not specified.
- If the bandwidth limit is decreased, no currently admitted OIFs are blocked. However, no new OIFs are admitted until the total admitted bandwidth for the interface drops below the new limit.
- If the bandwidth is increased to the point that the bandwidth limit for an interface is now exceeded, no currently admitted OIFs for the affected mroutes are blocked. However, no new OIFs are admitted until the total admitted bandwidth drops below the configured limit.

NOTE: **If the multicast bandwidth map that includes the set admission-bandwidth command is changed, all affected mroutes are reevaluated in the same manner described previously.**

As an example of this function, if the interface has accepted a total bandwidth of 2000000 bits per second (bps), and you set a limit of 1000000 bps on the interface, the router does not disconnect any already connected OIFs but prevents the interfaces from accepting any more groups. Over time, some groups leave the interfaces and, eventually, the interface limit of 1000000 bps is reached and maintained by the router.

If you set limits for both a port and interfaces on that port, the router uses the lower of the two limits when determining whether or not an interface can accept any new Multicast Listener Discovery (MLD) groups. For example, if you specify an admission bandwidth limit of 2000000 bps for the port and 3000000 bps groups for each interface, additional groups can only be accepted until the port limit of 2000000 bps is reached.

Related Documentation

- Before You Begin on page 164
- Example: Configuring an IPv6 Multicast Bandwidth Map on page 171
- Interface-Level Admission Bandwidth Limits for IPv6 Overview on page 182
- Enabling Interface-Level Admission Bandwidth Limitation for IPv6 on page 183
- OIF Interface Reevaluation for IPv6 on page 184

- ipv6 block-multicast-sources
- ipv6 mld version
- ipv6 multicast admission-bandwidth-limit
- set admission-bandwidth

Port-Level Multicast Traffic Configuration for IPv6

You can limit outgoing interface creation and limit the multicast bandwidth admitted on a port.

Creating IPv6 Mroute Port Limits

When a multicast forwarding entry (that is, an mroute) is added with an OIF on a port, the OIF count for that port is incremented. If you configure a port limit and the OIF count on the port count exceeds that limit, no OIFs on that port are added to mroutes (that is, new OIFs are blocked).

To configure a limit on the number of mroute OIFs that can be added across different virtual routers on a port:

- Issue the **mroute port limit** command in Global Configuration mode.

  host1(config)#**mroute port 3/0 limit 10**

Use the **no** version to remove any OIF port limits.

Port-Level Admission Bandwidth Limits for IPv6 Overview

Port-level multicast admission control is performed when an outgoing (OIF) on that port is added to the mroute for a given (S,G) multicast data stream and the multicast bandwidth map contains a **set admission-bandwidth** action for that (S,G).

When enabled, the admission-bandwidth for a particular (S,G) is read from the multicast bandwidth map and recorded in the mroute when the (S,G) mroute is created. When an IOF is subsequently added to the mroute, the OIF is blocked from forwarding data if the additional bandwidth contributed by the (S,G) would exceed the admission-bandwidth limit for the port on which the interface resides.

CAUTION: Before you can limit port-level admission bandwidth, you must first create a bandwidth map. See “Example: Configuring an IPv6 Multicast Bandwidth Map” on page 171 for details.

Enabling Port-Level Admission Bandwidth Limitation for IPv6

You can use the **mroute port admission-bandwidth-limit** command to limit the total multicast bandwidth that can be admitted on a port. The admitted bandwidth is summed across all virtual routers with IPv4 and IPv6 mroutes that have OIFs on the port.

NOTE: Admission bandwidth values for a given (S,G) mroute are determined from the bandwidth map. See "Example: Configuring an IPv6 Multicast Bandwidth Map" on page 171 for details.

To configure a limit on the admission bandwidth of OIFs containing IPv4 or IPv6 mroutes, across different virtual routers, on a port:

- Issue the **mroute port admission-bandwidth-limit** command in Global Configuration mode.

  ```
  host1(config)#mroute port admission-bandwidth-limit 3000000
  ```

Use the **no** version to remove any OIF admission bandwidth limits.

OIF Port Reevaluation for IPv6

If you change the admission bandwidth for a port, all mroutes with an outgoing interface (OIF) on that port are reevaluated as follows:

- If the bandwidth limit is increased, blocked OIFs can become unblocked. However, the order in which the mroutes are visited, and which (S,G) streams become unblocked, is not specified.
- If the bandwidth limit of a port is decreased, no currently admitted OIFs are blocked. However, no new OIFs are admitted until the total admitted bandwidth for the port drops below the new limit.
- If the bandwidth is increased to the point that the bandwidth limit for an interface is now exceeded, no currently admitted OIFs for the affected mroutes are blocked. However, no new OIFs are admitted until the total admitted bandwidth drops below the configured limit.

  NOTE: If the multicast bandwidth map that includes the **set admission-bandwidth** command is changed, all affected mroutes are reevaluated in the same manner described previously.

As an example of this function, if the port has accepted a total bandwidth of 3000000 bits per second (bps), and you set a limit of 2000000 bps on the port, the router does not disconnect any already connected OIFs but prevents the interfaces from accepting any more groups. Over time, some groups leave the interfaces and, eventually, the port limit of 2000000 bps is reached and maintained by the router.

If you set limits for both a port and interfaces on that port, the router uses the lower of the two limits when determining whether or not an interface can accept any new Multicast Listener Discovery (MLD) groups. For example, if you specify an admission bandwidth limit of 2000000 bps for the port and 3000000 bps groups for each interface, additional groups can only be accepted until the port limit of 2000000 bps is reached.

Related Documentation

- Before You Begin on page 164
- Port-Level Admission Bandwidth Limits for IPv6 Overview on page 185
- Enabling Port-Level Admission Bandwidth Limitation for IPv6 on page 185
- OIF Port Reevaluation for IPv6 on page 186
- mroute port admission-bandwidth-limit
- mroute port limit
- set admission-bandwidth

Deleting IPv6 Multicast Forwarding Entries

You can clear one or more forwarding entries from the multicast routing table. However, if you do so, the entries may reappear in the routing table if they are rediscovered. If you specify an *, the router clears all IP multicast forwarding entries. If you specify the IPv6 address of a multicast group, the router clears all multicast forward entries for that group. If you specify the IPv6 address of a multicast group and the IPv6 address of a multicast source, the router clears the multicast entry that matches that group and source.

To delete IPv6 multicast forwarding entries:

- Issue the **clear ipv6 mroute** command in Privileged Exec mode.

  host1:boston#**clear ipv6 mroute ***

Related Documentation

- Before You Begin on page 164
- Defining Permanent IPv6 Multicast Forwarding Entries on page 166
- clear ipv6 mroute

BGP Multicasting

BGP multicasting (MBGP) is an extension of the BGP unicast routing protocol. Many of the functions available for BGP unicasting are also available for MBGP.

The MBGP extensions specify that BGP can exchange information within different types of *address families*. The address families available are unicast IPv4, multicast IPv4, and VPN-IPv4. When you enable BGP, the router employs unicast IPv4 addresses by default.

We recommend you be thoroughly familiar with BGP before configuring MBGP. See *Configuring BGP Routing* in the *JunosE BGP and MPLS Configuration Guide* for detailed information about BGP and MBGP.

This topic discusses the following:

- Investigating Multicast Routes on page 188

Investigating Multicast Routes

You can use the **mtrace** command to trace the path that multicast packets take from a source to a destination through a multicast group address. This command is similar to the **traceroute** command for investigating unicast routes.

```
host1#mtrace 100.4.4.4 40.1.1.1 232.1.1.1
Tracing multicast route from 100.4.4.4 to 40.1.1.1 for group 232.1.1.1 using response address
  10.6.129.56
(Press ^c to stop.)
Received mtrace response packet of length 88
1. 40.1.1.1 Protocol: PIM(3) FwdingCode: RPF iif(9)
2. 21.2.2.2 Protocol: PIM(3) FwdingCode: Reached RP(8)
```

Related Documentation

- mtrace

CHAPTER 10

Monitoring IPv6 Multicast

IPv6 multicast enables a device to send packets to a group of hosts rather than to a list of individual hosts. This chapter describes how to monitor IPv6 multicast on the E Series router; it contains the following sections:

Monitoring Available IPv6 Routes for Reverse-Path Forwarding

Purpose Display the IPv6 routes that the router can use for RPF. You can specify the IPv6 address and the network mask to view routes to a particular destination. You can use the **detail** keyword to view more detailed information about routes to a particular destination. You can specify a unicast routing protocol to view routes associated with that protocol.

Action To display all IPv6 routes that the router can use for RPF:

```
host1#show ipv6 rpf-route
Protocol/Route type codes:
  O- OSPF, E1- external type 1, E2- external type2,
  N1- NSSA external type1, N2- NSSA external type2
  L- MPLS label, V- VR/VRF, *- indirect next-hop

         Prefix/Length             Type    Dst/Met         Intf
-------------------------------- --------- ------- --------------------
11:1:1:10::/60                   Static    1/0     ATM2/0.300
21:2:2:20::/60                   Static    1/0     ATM2/0.300
```

```
31:2:2:20::/60                     Connect    0/0      ATM2/0.300
131:1:1:10::/60                    Connect    0/0      ATM2/1.1300
1000::/64                          Static     1/0      ATM2/0.300
```

To display more detailed information about an IPv6 route to a particular destination:

`host1#`**show ipv6 rpf-route 1000::/64 detail**

```
1000::/64 Type:Static  Distance:1  Metric:0
        NextHop:31:2:2:23::2:3  IntfIndex 18  Intf ATM2/0.300
```

Meaning Table 51 on page 190 lists the **show ipv6 rpf-route** command output fields.

Table 51: show ipv6 rpf-route Output Fields

Field Name	Field Description
Protocol/Route type codes	Protocol and route type codes for the table that follows
Prefix	Value of the logical AND of the IPv6 address of the destination network and the subnet address
Length	Length of the subnet mask in bits
Type	• Connect—Subnet directly connected to the interface • Static—Static route
Dst	Distance configured for this route
Met	Learned or configured cost associated with this route
Intf	Type of interface and interface specifier for the next hop. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.

Related Documentation

- Defining IPv6 Static Routes for Reverse-Path Forwarding on page 165
- show ipv6 rpf-route

Monitoring IPv6 Hardware Multicast Packet Replication

You can use the commands explained in this topic to monitor hardware multicast packet replication-related statistics for port 8. For port 8, queue statistics have no direct relationship to any of the 8 ports because each packet that is transmitted through the queue can be sent through 1 or more of the 8 physical ports.

Multicast traffic redirected by the hardware multicast packet replication feature is displayed in the statistics for the IPv6 or VLAN interface over port 8, not the original IP or VLAN interface over the physical port.

The statistics for the IPv6 or VLAN interface over port 8 reflect the number of packets that passed through this interface destined for the hardware multicast packet replication hardware. These statistics have no direct correlation to the number of packets being transmitted from any of the physical ports.

The various statistics that are monitored over port 8 are:

- Monitoring Port Statistics on page 191
- Monitoring VLAN Statistics on page 196
- Monitoring IPv6 Statistics on page 198
- Monitoring MLD Statistics on page 212

Monitoring Port Statistics

Purpose Display the status of Gigabit Ethernet interfaces. You can use the **delta** keyword to view baselined statistics. You can use the **brief** keyword to view the operational status of all configured interfaces.

NOTE: The E120 and E320 routers output for **monitor** and **show** commands is identical to output from other E Series routers, except that the E120 and E320 router output also includes information about the adapter identifier in the interface specifier (*slot/adapter/port*).

Action To display the status of a Gigabit Ethernet interface:

```
host1#show interfaces gigabitEthernet 14/0/0
 GigabitEthernet14/0/0 is Up, Administrative status is Up
 Hardware is Intel IXF1104, address is 0090.1a42.0b87
 MAU is 1000BASE-SX
 TX Output Power: 469.6 uW RX Input Power: 0.5 uW
 MTU: Operational 1518, Administrative 1518
 Duplex Mode: Operational Full Duplex, Administrative Auto Negotiate
 Speed: Operational 1000 Mbps, Administrative Auto Negotiate
 Debounce: State is Disabled
 5 minute input rate 0 bits/sec, 0 packets/sec
 5 minute output rate 0 bits/sec, 0 packets/sec
 In: Bytes 0, Unicast 0
 Multicast 0, Broadcast 0
 Errors 0, Discards 0, Mac Errors 0, Alignment 0
 CRC 0, Too Longs 0, Symbol Errors 0
 Out: Bytes 0, Unicast 0
 Multicast 0, Broadcast 0
 Errors 0, Discards 0, Mac Errors 0, Deferred 0, No Carrier 0
 Collisions: Single 0, Multiple 0, Late 0, Excessive 0
 Policed Statistics:
 In: 0, Out: 0
 ARP Statistics:
 In: ARP requests 0, ARP responses 0
 Errors 0, Discards 0
 Out: ARP requests 0, ARP responses 0
```


Meaning Table 52 on page 192 lists the output fields of the **show interfaces gigabitEthernet** command.

Table 52: show interfaces gigabitEthernet Output Fields

Field Name	Field Description
GigabitEthernet or tenGigabitEthernet *interfaceSpecifier*	Status of the hardware on this interface: • up—Hardware is operational • down—Hardware is not operational
Administrative status	Operational state that you configured for this interface: • up—Interface is enabled • down—Interface is disabled
Hardware	Type of MAC device on this interface
Address	MAC address of the processor on this interface
MAU	Type of MAU on the primary and secondary physical ports: • SFP—1000BASE-LH, 1000BASE-SX, 1000BASE-ZX; for SFPs that are empty, SFP (Empty) appears in this field; for SFPs that are installed in the OC3-2 GE APS I/O module and do not have a Juniper Networks part number programmed, SFP (GE Compliant) appears in this field • XFP—10GBASE-SR (10 Gbps), 10GBASE-LR (10 Gbps), 10GBASE-ER (10 Gbps); for XFPs that are empty, XFP (Empty) appears in this field
TX Output Power	Transmitted output optical power
RX Input Power	Received input optical power
MTU	Size of the MTU for this interface: • Operational—Size of the largest packet processed • Administrative—Setting for MTU size that you specified
Duplex Mode	Duplex option for this interface: • Operational—Duplex option currently used • Administrative—Setting for duplex that you specified

Table 52: show interfaces gigabitEthernet Output Fields *(continued)*

Field Name	Field Description
Speed	Line speed for this interface: • Operational—Current rate at which packets are processed • Administrative—Setting for line speed that you specified
Debounce	Debounce configuration for this interface: • State is—Enabled, Disabled • Interval is—Number of seconds that this interface maintains a given state before the state change is reported to the upper-layer links
Clear arp	State of the removal of the ARP entries on an interface with redundant ports: • Enabled—Clears ARP entries on the interface when the primary link fails • Disabled—Maintains ARP entries on the interface until the specified timeout elapses

Link	Link information for this interface:
	• Operational Link Selected—Port that the I/O module is currently using: primary or secondary
	• Administrative link selected—Port that the I/O module is configured to use:
	  ◦ primary—Only primary port is configured to operate
	  ◦ secondary—Only redundant port is configured to operate
	  ◦ automatically—Software controls port redundancy automatically

Link Failover Timeout	Time to wait for a failed link to be active before the router uses a different active link
Primary link selected x times	Number of times that the I/O has used the primary port since the module was last rebooted
Secondary link selected x times	Number of times that the I/O has used the secondary port since the module was last rebooted`
Primary/Secondary link signal detected, Primary/Secondary link signal not detected	Specifies the port (primary or secondary) on which the router detects a signal (not displayed on GE I/O modules that do not support SFPs)
5 minute input rate	Data rates based on the traffic received in the last 5 minutes

Table 52: show interfaces gigabitEthernet Output Fields *(continued)*

Field Name	Field Description
5 minute output rate	Data rates based on the traffic sent in the last 5 minutes
In	Analysis of inbound traffic on this interface: • Bytes—Number of bytes received in error-free packets • Unicast—Number of unicast packets received • Multicast—Number of multicast packets received • Broadcast—Number of broadcast packets received • Errors—Total number of errors in all received packets; some packets might contain more than one error • Discards—Total number of discarded incoming packets • Mac Errors—Number of incoming packets discarded because of MAC sublayer failures • Alignment—Number of incomplete octets received • CRC—Number of packets discarded because the checksum that the router computed from the data does not match the checksum generated by the originating devices • Too Longs—Number of packets discarded because the size exceeded the MTU • Symbol Errors—Number of symbols received that the router did not correctly decode
Out	Analysis of outbound traffic on this interface: • Bytes—Number of bytes sent • Unicast—Number of unicast packets sent • Multicast—Number of multicast packets sent • Broadcast—Number of broadcast packets sent • Errors—Total number of errors in all transmitted packets; note that some packets might contain more than one error • Discards—Total number of discarded outgoing packets • Mac Errors—Number of outgoing packets discarded because of MAC sublayer failures • Deferred—Number of packets that the router delayed sending because the line was busy. In half duplex mode, a high number of deferrals means the link is very busy with traffic from other stations. In full duplex mode, when the link is always available for transmission, this number is zero. • No Carrier—Number of packets sent when carrier sense was unavailable

Table 52: show interfaces gigabitEthernet Output Fields *(continued)*

Field Name	Field Description
Collisions	Analysis of the collisions that occurred : • Single—Number of packets sent after one collision • Multiple—Number of packets sent after multiple collisions • Late—Number of packets aborted during sending because of collisions after 64 bytes • Excessive—Number of packets not sent because of too many collisions
Policed Statistics	Number of packets that exceeded the number allowed and were policed (or dropped)
ARP Statistics	Analysis of ARP traffic on this interface; In fields are for traffic received on the interface and Out fields are for traffic sent on the interface: • ARP requests—Number of ARP requests • ARP responses—Number of ARP responses • Errors—Total number of errors in all ARP packets • Discards—Total number of discarded ARP packets
Administrative qos-shaping-mode	QoS shaping mode: • disabled—Shaping mode is configured but not operational • frame—Statistics are reported about bytes in frames, such as transmitted bytes and dropped bytes. • cell—Shaping mode for shaping and policing rates is cell-based; resulting traffic stream conforms exactly to the policing rates configured in downstream devices. Reports statistics in bytes in cells and accounts for cell encapsulation and padding overhead. • none—Shaping mode is not configured
Operational qos-shaping-mode	Actual shaping mode for the interface: • disabled • frame • cell • none

Table 52: show interfaces gigabitEthernet Output Fields *(continued)*

Field Name	Field Description
queue	Hardware packet queue associated with the specified traffic class and interface : • traffic class—Name of traffic class • bound to—Interface to which queue is bound • Queue length—Length of the queue, in bytes • Forwarded packets, bytes—Number of packets and bytes that were forwarded on this queue • Dropped committed packets, bytes—Number of committed packets and bytes that were dropped • Dropped conformed packets, bytes—Number of conformed packets and bytes that were dropped • Dropped exceeded packets, bytes—Number of exceeded packets and bytes that were dropped

Monitoring VLAN Statistics

Purpose Display configuration and status information for a specified VLAN subinterface or for all VLAN subinterfaces configured on the router. You can use the **mac-address** keyword to display information about the VLAN subinterfaces that were configured with unique MAC addresses.

Action To display full status and configuration information for the specified VLAN subinterface:

```
host1#show vlan subinterface fastEthernet 0/0.1
    Interface       Status  MTU  Svlan Id Vlan Id Ethertype Type
------------------- ------ ---- -------- ------ --------- -------
FastEthernet 0/0.1    Up    1526    1      0    0x9100    Static

In: Bytes 39256, Packets 612
   Multicast 0, Broadcast 0
   Errors 0, Discards 0
  Out: Bytes 4538652, Packets 70911
   Multicast 0, Broadcast 70296
   Errors 0, Discards 0
ARP Statistics:
  In: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
  Out: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
```

Meaning Table 53 on page 196 lists the **show vlan subinterface** command output fields.

Table 53: show vlan subinterface Output Fields

Field Name	Field Description
Interface	Type and specifier of the VLAN subinterface
Status	Status of the VLAN subinterface: up, down, dormant, lowerLayerDown, absent

Table 53: show vlan subinterface Output Fields *(continued)*

Field Name	Field Description
MTU	Maximum allowable size (in bytes) of the MTU for the VLAN subinterface
Svlan Id	S-VLAN ID value, if configured
Vlan Id	VLAN ID value for the VLAN subinterface
Ethertype	S-VLAN Ethertype value, if configured
Type	Type of VLAN subinterface: • Static—VLAN or S-VLAN subinterface was configured statically • Dynamic—VLAN or S-VLAN subinterface was configured dynamically
In	Analysis of inbound traffic on this interface: • Bytes—Number of bytes received on the VLAN or S-VLAN subinterface • Packets—Sum of all unicast, broadcast, and multicast packets received on the VLAN or S-VLAN subinterface • Multicast—Number of multicast packets received on the VLAN or S-VLAN subinterface • Broadcast—Number of broadcast packets received on the VLAN or S-VLAN subinterface • Errors—Total number of errors in all received packets; some packets might contain more than one error • Discards—Total number of discarded incoming packets
Out	Analysis of outbound traffic on this interface: • Bytes—Number of bytes sent on the VLAN or S-VLAN subinterface • Packets—Number of packets sent on the VLAN or S-VLAN subinterface • Multicast—Number of multicast packets received on the VLAN or S-VLAN subinterface • Broadcast—Number of broadcast packets received on the VLAN or S-VLAN subinterface • Errors—Total number of errors in all transmitted packets; some packets might contain more than one error • Discards—Total number of discarded outgoing packets

Table 53: show vlan subinterface Output Fields *(continued)*

Field Name	Field Description
ARP Statistics	Analysis of ARP traffic on this interface; In fields are for traffic received on the interface and Out fields are for traffic sent on the interface: • ARP requests—Number of ARP requests • ARP responses—Number of ARP responses • Errors—Total number of errors in all ARP packets • Discards—Total number of discarded ARP packets

Monitoring IPv6 Statistics

Purpose Display detailed or summary information for a particular IPv6 interface or for all IPv6 interfaces.

Action To display statistics for all IPv6 interfaces:

```
host1#show ipv6 interface
null0 line protocol IpLoopback is up, ipv6 is up
  Network Protocols: IPv6
  Link local address: fe80::90:1a00:740:1d44
  Unnumbered Interface: Corresponding Numbered Interface not specified or
removed
  Operational MTU 1500 Administrative MTU 0
  Operational speed 100000000 Administrative speed 0
  Creation type Static
  Neighbor Discovery is disabled

  In Received Packets 0, Bytes 0
    Unicast Packets 0, Bytes 0
    Multicast Packets 0, Bytes 0
  In Total Dropped Packets 0, Bytes 0
    In Policed Packets 0
    In Invalid Source Address Packets 0
    In Error Packets 0
    In Discarded Packets 0

  Out Forwarded Packets 0, Bytes 0
    Unicast Packets 0, Bytes 0
    Multicast Routed Packets 0, Bytes 0
  Out Total Dropped Packets 0, Bytes 0
    Out Scheduler Dropped Packets 0, Bytes 0
    Out Policed Packets 0
    Out Discarded Packets 0

FastEthernet9/1.5 line protocol VlanSub is up, ipv6 is up
  Description: IPv6 interface in Virtual Router Hop5
  Network Protocols: IPv6
  Link local address: fe80::90:1a00:740:31ce
  Internet address: 5:1:1::2/64
  Operational MTU 1500 Administrative MTU 0
  Operational speed 100000000 Administrative speed 0
  Creation type Static
  ND reachable time is 3600000 milliseconds
  ND duplicate address detection attempts is 100
```



```
  ND neighbor solicitation retransmission interval is 1000 milliseconds
  ND proxy is enabled

  In Received Packets 13, Bytes 1356
    Unicast Packets 5, Bytes 588
    Multicast Packets 8, Bytes 768
  In Total Dropped Packets 0, Bytes 0
    In Policed Packets 0
    In Invalid Source Address Packets 0
    In Error Packets 0
    In Discarded Packets 0

  Out Forwarded Packets 22, Bytes 2480
    Unicast Packets 22, Bytes 2480
    Multicast Routed Packets 0, Bytes 0
  Out Total Dropped Packets 8, Bytes 0
    Out Scheduler Dropped Packets 0, Bytes 0
    Out Policed Packets 0
    Out Discarded Packets 8

  queue 0: traffic class best-effort, bound to ipv6 FastEthernet9/1.5
    Queue length 0 bytes
    Forwarded packets 4, bytes 680
    Dropped committed packets 0, bytes 0
    Dropped conformed packets 0, bytes 0
    Dropped exceeded packets 0, bytes 0

FastEthernet9/0.6 line protocol VlanSub is up, ipv6 is up
  Description: IPv6 interface in Virtual Router Hop6
  Network Protocols: IPv6
  Link local address: fe80::90:1a00:740:31cd
  Internet address: 6:1:1::1/64
  Operational MTU 1500 Administrative MTU 0
  Operational speed 100000000 Administrative speed 0
  Creation type Static
  ND reachable time is 3600000 milliseconds
  ND duplicate address detection attempts is 100
  ND neighbor solicitation retransmission interval is 1000 milliseconds
  ND proxy is enabled
ND RA source link layer is advertised
  ND RA interval is 200 seconds, lifetime is 1800 seconds
  ND RA managed flag is disabled, other config flag is disabled
  ND RA advertising prefixes configured on interface

  In Received Packets 0, Bytes 0
    Unicast Packets 0, Bytes 0
    Multicast Packets 0, Bytes 0
  In Total Dropped Packets 0, Bytes 0
    In Policed Packets 0
    In Invalid Source Address Packets 0
    In Error Packets 0
    In Discarded Packets 0

  Out Forwarded Packets 8, Bytes 768
    Unicast Packets 8, Bytes 768
    Multicast Routed Packets 0, Bytes 0
  Out Total Dropped Packets 5, Bytes 0
    Out Scheduler Dropped Packets 0, Bytes 0
    Out Policed Packets 0
    Out Discarded Packets 5

  queue 0: traffic class best-effort, bound to ipv6 FastEthernet9/0.6
    Queue length 0 bytes
    Forwarded packets 0, bytes 0
```

```
      Dropped committed packets 0, bytes 0
      Dropped conformed packets 0, bytes 0
      Dropped exceeded packets 0, bytes 0

  loopback5 line protocol IpLoopback is up, ipv6 is up
    Network Protocols: IPv6
    Link local address: fe80::90:1a00:740:1d44
    Internet address: 10:1:1:0:290:1aff:fe40:1d44/64 (eui-64)
    Operational MTU 1500 Administrative MTU 0
    Operational speed 100000000 Administrative speed 0
    Creation type Static
    Neighbor Discovery is disabled

    In Received Packets 0, Bytes 0
      Unicast Packets 0, Bytes 0
      Multicast Packets 0, Bytes 0
    In Total Dropped Packets 0, Bytes 0
      In Policed Packets 0
      In Invalid Source Address Packets 0
      In Error Packets 0
      In Discarded Packets 0

    Out Forwarded Packets 0, Bytes 0
      Unicast Packets 0, Bytes 0
      Multicast Routed Packets 0, Bytes 0
    Out Total Dropped Packets 0, Bytes 0
      Out Scheduler Dropped Packets 0, Bytes 0
      Out Policed Packets 0
      Out Discarded Packets 0

  IPv6 policy input ipv6InPol25
      rate-limit-profile Rlp2Mb classifier-group clgA entry 1
        Committed: 0 packets, 0 bytes
        Conformed: 0 packets, 0 bytes
        Exceeded:  0 packets, 0 bytes
      rate-limit-profile Rlp8Mb
        Committed: 0 packets, 0 bytes
        Conformed: 0 packets, 0 bytes
        Exceeded:  0 packets, 0 bytes
    IPv6 policy output ipv6PolOut2
      rate-limit-profile RlpOutA classifier-group clgB entry 1
        Committed: 0 packets, 0 bytes
        Conformed: 0 packets, 0 bytes
        Exceeded:  0 packets, 0 bytes
      rate-limit-profile RlpOutB
        Committed: 0 packets, 0 bytes
        Conformed: 0 packets, 0 bytes
        Exceeded:  0 packets, 0 bytes
    IPv6 policy local-input ipv6PolLocIn5
      rate-limit-profile Rlp1Mb classifier-group clgC entry 1
        Committed: 0 packets, 0 bytes
        Conformed: 0 packets, 0 bytes
        Exceeded:  0 packets, 0 bytes
      rate-limit-profile Rlp5Mb
        Committed: 0 packets, 0 bytes
        Conformed: 0 packets, 0 bytes
        Exceeded:  0 packets, 0 bytes
    queue 0: traffic class best-effort, bound to ipv6 FastEthernet9/0.6
      Queue length 0 bytes
      Forwarded packets 0, bytes 0
      Dropped committed packets 0, bytes 0
      Dropped conformed packets 0, bytes 0
      Dropped exceeded packets 0, bytes 0
```


To display statistics for a particular IPv6 interface:

```
host1#show ipv6 interface FastEthernet 9/0.6
FastEthernet9/0.6 line protocol VlanSub is up, ipv6 is up
  Description: IPv6 interface in Virtual Router Hop6
  Network Protocols: IPv6
  Link local address: fe80::90:1a00:740:31cd
  Internet address: 6:1:1::1/64
  Operational MTU 1500 Administrative MTU 0
  Operational speed 100000000 Administrative speed 0
  Creation type Static
  ND reachable time is 3600000 milliseconds
  ND duplicate address detection attempts is 100
  ND neighbor solicitation retransmission interval is 1000 milliseconds
  ND proxy is enabled
ND RA source link layer is advertised
  ND RA interval is 200 seconds, lifetime is 1800 seconds
  ND RA managed flag is disabled, other config flag is disabled
  ND RA advertising prefixes configured on interface

  In Received Packets 0, Bytes 0
    Unicast Packets 0, Bytes 0
    Multicast Packets 0, Bytes 0
  In Total Dropped Packets 0, Bytes 0
    In Policed Packets 0
    In Invalid Source Address Packets 0
    In Error Packets 0
    In Discarded Packets 0

  Out Forwarded Packets 8, Bytes 768
    Unicast Packets 8, Bytes 768
    Multicast Routed Packets 0, Bytes 0
  Out Total Dropped Packets 5, Bytes 0
    Out Scheduler Dropped Packets 0, Bytes 0
    Out Policed Packets 0
    Out Discarded Packets 5

  queue 0: traffic class best-effort, bound to ipv6 FastEthernet9/0.6
    Queue length 0 bytes
    Forwarded packets 0, bytes 0
    Dropped committed packets 0, bytes 0
    Dropped conformed packets 0, bytes 0
    Dropped exceeded packets 0, bytes 0

IPv6 policy input ipv6InPol25
    rate-limit-profile Rlp2Mb classifier-group clgA entry 1
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
    rate-limit-profile Rlp8Mb
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
  IPv6 policy output ipv6PolOut2
    rate-limit-profile RlpOutA classifier-group clgB entry 1
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
    rate-limit-profile RlpOutB
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
  IPv6 policy local-input ipv6PolLocIn5
```

```
    rate-limit-profile Rlp1Mb classifier-group clgC entry 1
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
    rate-limit-profile Rlp5Mb
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
  queue 0: traffic class best-effort, bound to ipv6 FastEthernet9/0.6
    Queue length 0 bytes
    Forwarded packets 0, bytes 0
    Dropped committed packets 0, bytes 0
    Dropped conformed packets 0, bytes 0
    Dropped exceeded packets 0, bytes 0
Http Redirect Url: http://www.juniper.net
```

To display detailed IPv6 status and configuration information for all IPv6 interfaces:

```
host1#show ipv6 interface detail
null0 line protocol IpLoopback is up, ipv6 is up
  Network Protocols: IPv6
  Link local address: fe80::90:1a00:740:1d44

  Unnumbered Interface: Corresponding Numbered Interface not specified or
removed
IPv6 statistics:
  Rcvd:  0 local destination
         0 hdr errors, 0 addr errors
         0 unkn proto, 0 discards
  Sent:  0 generated, 0 no routes, 0 discards

ICMPv6 statistics:
  Rcvd:  0 destination unreach, 0 admin unreach, 0 parameter problem
         0 time exceeded, 0 pkt too big, 0 echo requests
         0 echo replies
  Sent:  0 destination unreach, 0 admin unreach, 0 parameter problem
         0 time exceeded, 0 pkt too big, 0 echo requests
         0 echo replies

  Operational MTU 1500 Administrative MTU 0
  Operational speed 100000000 Administrative speed 0
  Creation type Static
  Neighbor Discovery is disabled

  ICMPv6 statistics:
    Rcvd:  0 total, 0 errors
           0 rtr solicits, 0 rtr advertisements
           0 neighbor solicits, 0 neighbor advertisements
           Group membership: 0 queries, 0 responses, 0 reductions
           0 redirects
    Sent:  0 total, 0 errors
           0 rtr solicits, 0 rtr advertisements
           0 neighbor solicits, 0 neighbor advertisements
           Group membership: 0 queries, 0 responses, 0 reductions
           0 redirects

  In Received Packets 0, Bytes 0
    Unicast Packets 0, Bytes 0
    Multicast Packets 0, Bytes 0
  In Total Dropped Packets 0, Bytes 0
    In Policed Packets 0
    In Invalid Source Address Packets 0
```



```
    In Error Packets 0
    In Discarded Packets 0

  Out Forwarded Packets 0, Bytes 0
    Unicast Packets 0, Bytes 0
    Multicast Routed Packets 0, Bytes 0
  Out Total Dropped Packets 0, Bytes 0
    Out Scheduler Dropped Packets 0, Bytes 0
    Out Policed Packets 0
    Out Discarded Packets 0

FastEthernet9/1.5 line protocol VlanSub is up, ipv6 is up
  Description: IPv6 interface in Virtual Router Hop5
  Network Protocols: IPv6
  Link local address: fe80::90:1a00:740:31ce

  Internet address: 5:1:1::2/64
IPv6 statistics:
  Rcvd:  0 local destination
         0 hdr errors, 0 addr errors
         0 unkn proto, 0 discards
  Sent:  0 generated, 0 no routes, 0 discards

ICMPv6 statistics:
  Rcvd:  0 destination unreach, 0 admin unreach, 0 parameter problem
         0 time exceeded, 0 pkt too big, 0 echo requests
         3 echo replies
  Sent:  0 destination unreach, 0 admin unreach, 0 parameter problem
         0 time exceeded, 0 pkt too big, 5 echo requests
         0 echo replies

  Operational MTU 1500 Administrative MTU 0
  Operational speed 100000000 Administrative speed 0
  Creation type Static
  ND reachable time is 3600000 milliseconds
  ND duplicate address detection attempts is 100
  ND neighbor solicitation retransmission interval is 1000 milliseconds
  ND proxy is enabled
ND RA source link layer is advertised
  ND RA interval is 200 seconds, lifetime is 1800 seconds
  ND RA managed flag is disabled, other config flag is disabled
  ND RA advertising prefixes configured on interface

  ICMPv6 statistics:
    Rcvd:  13 total, 0 errors
           0 rtr solicits, 8 rtr advertisements
           1 neighbor solicits, 1 neighbor advertisements
           Group membership: 0 queries, 0 responses, 0 reductions
           0 redirects
    Sent:  31 total, 0 errors
           0 rtr solicits, 16 rtr advertisements
           5 neighbor solicits, 5 neighbor advertisements
           Group membership: 0 queries, 0 responses, 0 reductions
           0 redirects

  In Received Packets 13, Bytes 1356
    Unicast Packets 5, Bytes 588
    Multicast Packets 8, Bytes 768
  In Total Dropped Packets 0, Bytes 0
    In Policed Packets 0
    In Invalid Source Address Packets 0
    In Error Packets 0
    In Discarded Packets 0
```

```
  Out Forwarded Packets 22, Bytes 2480
    Unicast Packets 22, Bytes 2480
    Multicast Routed Packets 0, Bytes 0
  Out Total Dropped Packets 8, Bytes 0
    Out Scheduler Dropped Packets 0, Bytes 0
    Out Policed Packets 0
    Out Discarded Packets 8

  queue 0: traffic class best-effort, bound to ipv6 FastEthernet9/1.5
    Queue length 0 bytes
    Forwarded packets 4, bytes 680
    Dropped committed packets 0, bytes 0
    Dropped conformed packets 0, bytes 0
    Dropped exceeded packets 0, bytes 0

FastEthernet9/0.6 line protocol VlanSub is up, ipv6 is up
  Description: IPv6 interface in Virtual Router Hop6
  Network Protocols: IPv6
  Link local address: fe80::90:1a00:740:31cd

  Internet address: 6:1:1::1/64
IPv6 statistics:
  Rcvd:  0 local destination
         0 hdr errors, 0 addr errors
         0 unkn proto, 0 discards
  Sent:  0 generated, 0 no routes, 0 discards

ICMPv6 statistics:
  Rcvd:  0 destination unreach, 0 admin unreach, 0 parameter problem
         0 time exceeded, 0 pkt too big, 0 echo requests
         0 echo replies
  Sent:  0 destination unreach, 0 admin unreach, 0 parameter problem
         0 time exceeded, 0 pkt too big, 0 echo requests
         0 echo replies

  Operational MTU 1500 Administrative MTU 0
  Operational speed 100000000 Administrative speed 0
  Creation type Static
  ND reachable time is 3600000 milliseconds
  ND duplicate address detection attempts is 100
  ND neighbor solicitation retransmission interval is 1000 milliseconds
  ND proxy is enabled
ND RA source link layer is advertised
  ND RA interval is 200 seconds, lifetime is 1800 seconds
  ND RA managed flag is disabled, other config flag is disabled
  ND RA advertising prefixes configured on interface

  ICMPv6 statistics:
    Rcvd:  0 total, 0 errors
           0 rtr solicits, 0 rtr advertisements
           0 neighbor solicits, 0 neighbor advertisements
           Group membership: 0 queries, 0 responses, 0 reductions
           0 redirects
    Sent:  13 total, 0 errors
           0 rtr solicits, 9 rtr advertisements
           2 neighbor solicits, 2 neighbor advertisements
           Group membership: 0 queries, 0 responses, 0 reductions
           0 redirects

  In Received Packets 0, Bytes 0
    Unicast Packets 0, Bytes 0
    Multicast Packets 0, Bytes 0
  In Total Dropped Packets 0, Bytes 0
    In Policed Packets 0
```



```
    In Invalid Source Address Packets 0
    In Error Packets 0
    In Discarded Packets 0

  Out Forwarded Packets 8, Bytes 768
    Unicast Packets 8, Bytes 768
    Multicast Routed Packets 0, Bytes 0
  Out Total Dropped Packets 5, Bytes 0
    Out Scheduler Dropped Packets 0, Bytes 0
    Out Policed Packets 0
    Out Discarded Packets 5

  queue 0: traffic class best-effort, bound to ipv6 FastEthernet9/0.6
    Queue length 0 bytes
    Forwarded packets 0, bytes 0
    Dropped committed packets 0, bytes 0
    Dropped conformed packets 0, bytes 0
    Dropped exceeded packets 0, bytes 0

loopback5 line protocol IpLoopback is up, ipv6 is up
  Network Protocols: IPv6
  Link local address: fe80::90:1a00:740:1d44

  Internet address: 10:1:1:0:290:1aff:fe40:1d44/64 (eui-64)
IPv6 statistics:
  Rcvd:  0 local destination
         0 hdr errors, 0 addr errors
         0 unkn proto, 0 discards
  Sent:  0 generated, 0 no routes, 0 discards

ICMPv6 statistics:
  Rcvd:  0 local destination
         0 hdr errors, 0 addr errors
         0 unkn proto, 0 discards
  Sent:  0 generated, 0 no routes, 0 discards

ICMPv6 statistics:
  Rcvd:  0 destination unreach, 0 admin unreach, 0 parameter problem
         0 time exceeded, 0 pkt too big, 0 echo requests
         0 echo replies
  Sent:  0 destination unreach, 0 admin unreach, 0 parameter problem
         0 time exceeded, 0 pkt too big, 0 echo requests
         0 echo replies

  Operational MTU 1500 Administrative MTU 0
  Operational speed 100000000 Administrative speed 0
  Creation type Static
  Neighbor Discovery is disabled

  ICMPv6 statistics:
    Rcvd:  0 total, 0 errors
           0 rtr solicits, 0 rtr advertisements
           0 neighbor solicits, 0 neighbor advertisements
           Group membership: 0 queries, 0 responses, 0 reductions
           0 redirects
    Sent:  0 total, 0 errors
           0 rtr solicits, 0 rtr advertisements
           0 neighbor solicits, 0 neighbor advertisements
           Group membership: 0 queries, 0 responses, 0 reductions
           0 redirects

  In Received Packets 0, Bytes 0
    Unicast Packets 0, Bytes 0
    Multicast Packets 0, Bytes 0
```

```
  In Total Dropped Packets 0, Bytes 0
    In Policed Packets 0
    In Invalid Source Address Packets 0
    In Error Packets 0
    In Discarded Packets 0

  Out Forwarded Packets 0, Bytes 0
    Unicast Packets 0, Bytes 0
    Multicast Routed Packets 0, Bytes 0
  Out Total Dropped Packets 0, Bytes 0
    Out Scheduler Dropped Packets 0, Bytes 0
    Out Policed Packets 0
    Out Discarded Packets 0
IPv6 policy input ipv6InPol25
    rate-limit-profile Rlp2Mb classifier-group clgA entry 1
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
    rate-limit-profile Rlp8Mb
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
  IPv6 policy output ipv6PolOut2
    rate-limit-profile RlpOutA classifier-group clgB entry 1
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
    rate-limit-profile RlpOutB
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
  IPv6 policy local-input ipv6PolLocIn5
    rate-limit-profile Rlp1Mb classifier-group clgC entry 1
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
    rate-limit-profile Rlp5Mb
      Committed: 0 packets, 0 bytes
      Conformed: 0 packets, 0 bytes
      Exceeded:  0 packets, 0 bytes
  queue 0: traffic class best-effort, bound to ipv6 FastEthernet9/0.6
    Queue length 0 bytes
    Forwarded packets 0, bytes 0
    Dropped committed packets 0, bytes 0
    Dropped conformed packets 0, bytes 0
    Dropped exceeded packets 0, bytes 0
```

To display brief summary of IPv6 status and configuration information for all IPv6 interfaces:

host1# **show ipv6 interface brief**

```
    Interface          IPv6-Address     Status  Protocol   Description
----------------- ------------------ ------- -------- ----------------
null0              Unnumbered         up      up
FastEthernet9/1.5  5:1:1::2/64        up      up      IPv6 interface in
                                                      Virtual Router Hop
                                                            5
FastEthernet9/0.6  6:1:1::1/64        up      up      IPv6 interface in
                                                      Virtual Router Hop
                                                             6
```



```
loopback5          10:1:1:0:290:1aff:fe up      up
                   40:1d44/64
```

Meaning Table 54 on page 207 lists the **show ipv6 interface** command output fields.

Table 54: show ipv6 interface Output Fields

Field Name	Field Description
Description	Optional description for the interface or address specified
Network Protocols	Network protocols configured on this interface
Link local address	Local IPv6 address of this interface
Internet address	External address of this interface
IPv6 statistics Rcvd	• local destination—Frames with this router as their destination • hdr errors—Number of packets containing header errors • addr errors—Number of packets containing addressing errors • unkn proto—Number of packets received containing unknown protocols • discards—Number of discarded packets
IPv6 statistics Sent	• generated—Number of packets generated • no routes—Number of packets that could not be routed • discards—Number of packets that could not be routed that were discarded NOTE: If you configure the router to discard packets for static routes with null 0 interfaces as the next-hop points using the **reject** keyword with the **ipv6 route** command, the value displayed in this field also includes the packets that reached the null 0 interface and were dropped.

Table 54: show ipv6 interface Output Fields *(continued)*

Field Name	Field Description
ICMPv6 statistics Rcvd	• total—Total number of received packets • errors—Error packets received • destination unreach—Packets received with destination unreachable • admin unreach—Packets received because the destination was administratively unreachable (for example, the packet encountered a firewall filter) • parameter problem—Packets received with parameter errors • time exceeded—Packets received with time-to-live exceeded • pkt too big—Number of packet-too-big messages received that indicate a packet was too large to forward because of the allowed MTU size • redirects—Received packet redirects • echo requests—Echo request (ping) packets • echo replies—Echo replies received • rtr solicits—Number of received router solicitations • rtr advertisements—Number of received router advertisements • neighbor solicits—Number of received neighbor solicitations • neighbor advertisements—Number or received neighbor advertisements • Group membership (queries, responses, reductions)—Number of queries, responses, and reduction requests received from within a group to which the interface is assigned

Table 54: show ipv6 interface Output Fields *(continued)*

Field Name	Field Description
ICMPv6 statistics Sent	• total—Total number of sent packets • errors—Error packets sent • destination unreach—Packets sent with destination unreachable NOTE: If you configure the router to discard packets for static routes with null 0 interfaces as the next-hop points using the **reject** keyword with the **ipv6 route** command, the value displayed in this field also includes the number of ICMPv6 unreachable messages sent out for packets that reached null 0 interfaces with static routes. • admin unreach—Packets sent because the destination was administratively unreachable (for example, due to a firewall filter) • parameter problem—Packets sent with parameter errors • time exceeded—Packets sent with time-to-live exceeded • pkt too big—Number of packet-too-big messages sent because a received packet was too large to forward because of the allowed MTU size • redirects—Sent packet redirects • echo requests—Echo request (ping) packets • echo replies—Echo replies sent • rtr solicits—Number of sent router solicitations • rtr advertisements—Number of sent router advertisements • neighbor solicits—Number of sent neighbor solicitations • neighbor advertisements—Number of sent neighbor advertisements • Group membership (queries, responses, reductions)—Number of queries, responses, and reduction requests sent to a group of which the interface is assigned
Operational MTU	Value of the MTU
Administrative MTU	Value of the MTU if it has been administratively overridden using the configuration
Operational speed	Speed of the interface
Administrative speed	Value of the speed if it has been administratively overridden using the configuration
Creation type	Method by which the interface was created (static or dynamic)

Table 54: show ipv6 interface Output Fields *(continued)*

Field Name	Field Description
HTTP Redirect Url	URL to which a subscriber's initial web browser session is redirected
ND reachable time	Amount of time (in milliseconds) that the neighbor is expected to remain reachable
ND duplicate address detection attempts	Number of times that the router attempts to determine a duplicate address
ND neighbor solicitation retransmission interval	Interval in which the router retransmits neighbor solicitations
ND proxy	Indicates whether the router will reply to solicitations on behalf of a known neighbor
ND RA source link layer	Indicates whether the RA includes the link layer
ND RA interval	Interval (in seconds) of the neighbor discovery router advertisement
ND RA lifetime	Lifetime (in seconds) of the neighbor discovery router advertisement
ND RA managed flag	State of the neighbor discovery router advertisement managed flag
ND RA other config flag	State of the neighbor discovery router advertisement other config flag
ND RA advertising prefixes	Configured advertisement prefixes for neighbor discovery router advertisement
In Received Packets, Bytes	Total number of packets and bytes received on this interface: • Unicast Packets, Bytes—Unicast packets and bytes received on the IPv6 interface; link-local received multicast packets (non-multicast-routed frames) are counted as unicast packets • Multicast Packets, Bytes—Multicast packets and bytes received on the IPv6 interface which are then multicast-routed are counted as multicast packets

Table 54: show ipv6 interface Output Fields *(continued)*

Field Name	Field Description
In Total Dropped Packets, Bytes	Total number of inbound packets and bytes dropped on this interface: • In Policed Packets—Packets that were received and dropped on the interface for any of the following reasons: exceeding the token bucket limit, exceeding the rate limit, a drop action in a policy, discarded MAC validation packets, a destination address lookup failure, or when the destination address is an IP interface that has a route configured to the null 0 interface. • In Invalid Source Address Packets—Packets received with invalid source address (for example, spoofed packets) • In Error Packets—Number of packets received with errors • In Discarded Packets—Packets received that were discarded for reasons other than rate limits, errors, and invalid source address
Out Forwarded Packets, Bytes	Total number of packets and bytes that were sent from this interface: • Unicast Packets, Bytes—Unicast packets and bytes that were sent from this interface • Multicast Routed Packets, Bytes—Multicast packets and bytes that were sent from this interface
Out Total Dropped Packets	Total number of outbound packets and bytes dropped by this interface: • Out Scheduler Dropped Packets, Bytes—Number of outbound packets and bytes dropped by the scheduler • Out Policed Packets, Bytes—Number of outbound packets and bytes dropped because of rate limits • Out Discarded Packets—Number of outbound packets that were discarded for reasons other than those dropped by the scheduler and those dropped because of rate limits
IPv6 policy	Type (input, output, local-input) and name of policy: • rate-limit-profile—Name of profile • classifier-group entry—Entry index • Committed—Number of packets and bytes conforming to the committed access rate • Conformed—Number of packets and bytes that exceed the committed access rate but conform to the peak access rate • Exceeded—Number of packets and bytes exceeding the peak access rate

Table 54: show ipv6 interface Output Fields *(continued)*

Field Name	Field Description
queue, traffic class, bound to ipv6	Queue and traffic class bound to the specified IPv6 interface: • Queue length—Number of bytes in queue • Dropped committed packets, bytes—Total number of committed packets and bytes dropped by this interface • Dropped conformed packets, bytes—Total number of conformed packets and bytes dropped by this interface • Dropped exceeded packets, bytes—Total number of exceeded packets and bytes dropped by this interface

Monitoring MLD Statistics

Purpose Display MLD information for interfaces on which you enabled MLD.

Action To display statistics, including hardware multicast packet replication status, for all IPv6 interfaces on which you enabled MLD:

```
host1:boston#show ipv6 mld interface

Interface ATM2/0.103 address 13.0.0.1/255.255.255.0
Administrative state enabled, Operational state enabled
Interface parameters:
    Version 2
    State Querier
    Query Interval 125 secs, 125 secs before the next query
    Other querier present interval 250 secs
    Maximum response time 255 (in 10ths of a second)
    Last member query interval 10 (in 10ths of a second)
    Robustness 3
    No inbound access group
    No inbound access source-group
    No inbound apply-oif-map
    Immediate Leave: disabled
Explicit Host Tracking: enabled
Max-Group limit: No Limit
Admission-Bandwidth limit: No limit
IOA Packet Replication: None
    Group Count: 1
Interface statistics:
    Rcvd: 2 reports, 0 leaves, 0 wrong version queries
    Sent: 2 queries
    Groups learnt: 1

Counts: 0 down, 0 init state, 1 querier, 0 non-querier, 1 Total
```

To display statistics, including hardware multicast packet replication status, for a particular IPv6 interface on which you enabled MLD:

```
host1#show ipv6 mld interface gigabitEthernet 3/0.0
Interface GigabitEthernet3/0.0 address 10.1.1.1/255.255.255.0
```



```
Administrative state enabled, Operational state enabled
Interface parameters:
    Version 1
    State Querier
    Query Interval 125 secs, 123 secs before the next query
    Other querier present interval 250 secs
    Maximum response time 100 (in 10ths of a second)
    Last member query interval 10 (in 10ths of a second)
    Robustness 3
    No inbound access group
    No inbound access source-group
    No inbound apply-oif-map
    Immediate Leave: disabled
Explicit Host Tracking: enabled
Max-Group limit: No Limit
    Group Count: 0
    IOA packet replication gigabitEthernet 3/8.1
Interface statistics:
    Rcvd: 0 reports, 0 leaves, 0 wrong version queries
    Sent: 14 queries
    Groups learnt: 0

Counts: 0 down, 0 init state, 1 querier, 0 non-querier, 1 Total
```

Meaning Table 55 on page 213 lists the **show ipv6 mld interface** command output fields.

Table 55: show ipv6 mld interface Output Fields

Field Name	Field Description
Interface	Type of interface and interface specifier. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.
address	IPv6 link-local address of the interface
Administrative state	Status of the interface in the software: enabled or disabled
Operational state	Physical status of the interface: enabled or disabled
Version	MLD version
State	Function of the interface: querier or nonquerier
Query Interval	Time interval at which this interface sends query messages
Other querier present interval	Time that the interface waits before declaring itself as the querier
Maximum response time	Time interval during which this interface expects a host to respond
Graceful restart	Status of graceful restart: active or complete

Table 55: show ipv6 mld interface Output Fields *(continued)*

Field Name	Field Description
Last member query interval	Time that this interface waits before sending a new query to a host that sends a group leave message
Robustness	Number of times this interface sends MLD messages
Inbound access group	Information about IPv6 access lists configured with the **ipv6 mld access-group** command
No inbound access group	No IPv6 access list configured with the **ipv6 mld access-group** command
Inbound access source-group	Information about IPv6 access lists configured with the **ipv6 mld access-source-group** command
No inbound access source-group	No IPv6 access list configured with the **ipv6 mld access-source-group** command
Inbound apply-oif-map	Information about OIF map names configured with the **ipv6 mld apply-oif-map** command
No inbound apply-oif-map	No OIF map name configured with the **ipv6 mld apply-oif-map** command
Immediate Leave	Setting of the **ipv6 mld immediate-leave** command: enabled or disabled
Explicit Host Tracking	Setting of the **ipv6 mld explicit-tracking** command: enabled or disabled
Max-Group limit	Number of MLD groups that the interface can accept, as configured with the **ipv6 mld group limit** command
Group Count	Number of MLD groups that the interface has accepted
IOA packet replication	Hardware multicast packet replication interface to which egress multicast packets on this interface are redirected

Table 55: show ipv6 mld interface Output Fields *(continued)*

Field Name	Field Description
Interface statistics Rcvd	Information about MLD messages received on this interface: • reports—Number of group multicast listener reports received • leaves—Number of group multicast listener done messages received • wrong version queries—Number of multicast listener queries received from devices running a different version of MLD • Interface statistics Sent—Number of MLD messages this interface has sent • Interface statistics Groups learnt—Number of groups this interface has discovered
Counts	Total number of MLD interfaces

Related Documentation

- Configuring IPv6 Hardware Multicast Packet Replication on page 181
- show interfaces
- show ipv6 interface
- show ipv6 mld interface
- show vlan subinterface

Monitoring IPv6 Multicast Forwarding Entries

Purpose Displays information about all or specified multicast forwarding entries. You can specify a multicast group IPv6 address or both a multicast group IPv6 address and a multicast source IPv6 address to display information about particular multicast forwarding entries.

Action To display all multicast forwarding entries while bandwidth rate is constant:

```
host1#show ipv6 mroute
IPv6 Multicast Routing Table

(S, G) uptime d h:m:s[, expires d h:m:s]
  [Admission bandwidth: bps]
  [QoS bandwidth: bps]
  RPF route: addr/mask, incoming interface
             neighbor address, owner route-owner
  Incoming interface list:
      Interface (addr/mask), State/Owner [(RPF IIF)]
  Outgoing interface list:
      Interface (addr/mask), State/Owner, Uptime/Expires

(10:0:0:1:1::, ff0e::1) uptime 0 01:04:12
  RPF route: 10:0:0:1::/64, incoming interface ATM2/3.1001
             neighbor 10:0:0:1::1, owner Local
  Incoming interface list:
      ATM2/3.1001 (10:0:0:1::1/64), Accept/Pim (RPF IIF)
```

```
  Outgoing interface list:
      ATM2/0.200 (21:2:2:21::2:1/60), Forward/Pim, 0 01:04:12/never

(10:0:0:1:2::, ff0e::1) uptime 0 01:04:12
  RPF route: 10:0:0:1::/64, incoming interface ATM2/3.1001
             neighbor 10:0:0:1::1, owner Local
  Incoming interface list:
      ATM2/3.1001 (10:0:0:1::1/64), Accept/Pim (RPF IIF)
  Outgoing interface list:
      ATM2/0.200 (21:2:2:21::2:1/60), Forward/Pim, 0 01:04:12/never

Counts: 2 (S, G) entries
        0 (*, G) entries
```

To display all multicast forwarding entries when bandwidth limit of all the join interfaces exceeds configured CAC limits:

```
host1#show ipv6 mroute
IP Multicast Routing Table

(S, G) uptime d h:m:s[, expires d h:m:s]
  [Admission bandwidth: bps]
  [QoS bandwidth: bps]
  RPF route: addr/mask, incoming interface
             neighbor address, owner route-owner
  Incoming interface list:
      Interface (addr/mask), State/Owner [(RPF IIF)]
  Outgoing interface list:
      Interface (addr/mask), State/Owner, Uptime/Expires

(10:0:0:1:1::, ff0e::1) uptime 0 01:04:12
  RPF route: 10:0:0:1::/64, incoming interface ATM2/3.1001
             neighbor 10:0:0:1::1, owner Local
  Incoming interface list:
      ATM2/3.1001 (10:0:0:1::1/64), Accept/Pim (RPF IIF)
  Outgoing interface list:
      ATM2/0.200 (21:2:2:21::2:1/60), Forward/Pim, 0 01:04:12/never

(10:0:0:1:2::, ff0e::1) uptime 0 01:04:12
  RPF route: 10:0:0:1::/64, incoming interface ATM2/3.1001
             neighbor 10:0:0:1::1, owner Local
  Incoming interface list:
      ATM2/3.1001 (10:0:0:1::1/64), Accept/Pim (RPF IIF)
  Outgoing interface list:
      ATM2/0.200 (21:2:2:21::2:1/60), Blocked (join-intf-adm-limit)/IGMP, 0
01:04:12/never

Counts: 2 (S, G) entries
        0 (*, G) entries
```

Meaning Table 56 on page 216 lists the **show ipv6 mroute** command output fields.

Table 56: show ipv6 mroute Output Fields

Field Name	Field Description
(S,G)	IPv6 addresses of the multicast source and the multicast group
Uptime	Length of time that the (S,G) pair has been active, in *days hours:minutes:seconds* format

Table 56: show ipv6 mroute Output Fields *(continued)*

Field Name	Field Description
Expires	Length of time for which the (S,G) pair will be active, in *days hours:minutes:seconds* format
Admission bandwidth	Admission bandwidth (in bps)
QoS bandwidth	QoS bandwidth (in bps)
RPF Route	IPv6 address and prefix of the RPF route
Incoming interface	Type and specifier of the incoming interface for the RPF route
neighbor address	IPv6 address of the neighbor
owner	Owner of the route: • Local—route belonging to the local interface • Static—Static route • Other protocols—Route established by a protocol
Incoming interface list	List of incoming interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: accept or discard • Multicast protocol that owns the interface • Time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Time that the interface will cease to be active in this multicast forwarding entry, in *days hours:minutes:seconds* format
Outgoing interface list	List of outgoing interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Forward or Blocked (intf-adm-limit, join-intf-adm-limit, port-adm-limit, port-limit, port-priority-limit) • Protocol running on the interface: PIM or MLD • Time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Time that the interface will cease to be active in this multicast forwarding entry, in *days hours:minutes:seconds* format

Table 56: show ipv6 mroute Output Fields *(continued)*

Field Name	Field Description
Counts	Numbers of types of source group mappings: • (S,G)—Number of (S,G) entries • (*,G)—Number of (*,G) entries

Related Documentation

- Defining IPv6 Static Routes for Reverse-Path Forwarding on page 165
- show ipv6 mroute

Monitoring Active IPv6 Multicast Routes

Purpose Display the active multicast routes.

- You can specify a multicast group IPv6 address or both a multicast group IPv6 address and a multicast source IPv6 address to display information about particular multicast forwarding entries.
- You can specify the bandwidth threshold to display the active multicast routes with admission bandwidth greater than the specified bandwidth threshold. The default bandwidth threshold is 4000 bps.
- You can use the **summary** option to see a summary rather than a detailed description.
- You can use the **count** option to display the number of active multicast forwarding entries.
- You can use the **oif-detail** option to display the details of the join interfaces corresponding to the mapped interface when oif-mapping is configured.
- You can use the **statistics** option to display statistics for packets received through all active multicast forwarding entries that the router has added to the multicast routing table and established on the appropriate line modules.

Action To display the active multicast routes with admission bandwidth above 10000 bps:

```
host1#show ipv6 mroute active 10000
Active IP Multicast Routes >=10000 bps

(S, G) uptime d h:m:s[, expires d h:m:s]

[Admission bandwidth: bps]

[QoS bandwidth: bps]

RPF route: addr/mask, incoming interface

neighbor address, owner route-owner

Incoming interface list:

Interface (addr/mask), State/Owner [(RPF IIF)]

Outgoing interface list:

Interface (addr/mask), State/Owner, Uptime/Expires
```



```
(52::1, ff3e::1) uptime 0 00:01:07
Admission bandwidth: 47000 bps (adaptive)
QoS bandwidth: 47000 bps (adaptive)
RPF route: 52::/112, incoming interface ATM2/1.17
neighbor 17::2, owner NetmgmtRpf
Incoming interface list:
ATM2/1.17 (fe80::90:1a00:3140:1ff8/128), Accept/MLD (RPF IIF)
Outgoing interface list:
NULL
Counts: 1 (S, G) entries
0 (*, G) entries
```

To display the summary of active multicast routes:

```
host1#show ipv6 mroute summary active
Active IP Multicast Routes >=4000 bps
Group Address Source Address  RPF route RPF Iif   #Oifs
--------------- ------------ ----------- --------- -----
232.0.0.1     51.0.0.1      51.0.0.0/24  ATM3/1.17  0
232.0.0.2     51.0.0.1      51.0.0.0/24  ATM3/1.17  0
232.0.0.3     51.0.0.1      51.0.0.0/24  ATM3/1.17  0
Counts: 3 (S, G) entries
0 (*, G) entries
```

Meaning Table 57 on page 219 lists the **show ipv6 mroute active** and **show ipv6 mroute summary active** commands output fields.

Table 57: show ipv6 mroute active and show ipv6 mroute summary active Output Fields

Field Name	Field Description
(S,G)	IPv6 addresses of the multicast source and the multicast group
Uptime	Length of time that the (S,G) pair has been active, in *days hours:minutes:seconds* format
Expires	Length of time for which the (S,G) pair will be active, in *days hours:minutes:seconds* format
Admission bandwidth	Admission bandwidth (in bps)
QoS bandwidth	QoS bandwidth (in bps)
RPF Route	IPv6 address and prefix of the RPF route

Table 57: show ipv6 mroute active and show ipv6 mroute summary active Output Fields *(continued)*

Field Name	Field Description
Incoming interface	Type and specifier of the incoming interface for the RPF route
neighbor address	IPv6 address of the neighbor
owner	Owner of the route: • Local—route belonging to the local interface • Static—Static route • Other protocols—Route established by a protocol
Incoming interface list	List of incoming interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: accept or discard • Multicast protocol that owns the interface • Time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Time that the interface will cease to be active in this multicast forwarding entry, in *days hours:minutes:seconds* format
Outgoing interface list	List of outgoing interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Forward or Blocked (intf-adm-limit, join-intf-adm-limit, port-adm-limit, port-limit, port-priority-limit) • Protocol running on the interface: PIM or MLD • Time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Time that the interface will cease to be active in this multicast forwarding entry, in *days hours:minutes:seconds* format
Counts	Numbers of types of source group mappings: • (S,G)—Number of (S,G) entries • (*,G)—Number of (*,G) entries
Group Address	IP address of the multicast group
Source Address	IP address of the multicast source

Table 57: show ipv6 mroute active and show ipv6 mroute summary active Output Fields *(continued)*

Field Name	Field Description
RPF Iif	Type and identifier for the incoming interface for the RPF route
#Oifs	Number of outgoing interfaces

Related Documentation

- Defining IPv6 Static Routes for Reverse-Path Forwarding on page 165
- show ipv6 mroute

Monitoring Multicast Routes on Virtual Router Ports

Purpose Display information for multicast routes on a port across all virtual routers.

NOTE: This command displays information for mroutes on a port across all virtual routers.

Action To display the multicast route port outgoing interface, limits, counts, bandwidth settings, and bandwidth accepted:

host1#**show mroute port count**

```
BW     Priority
Port    Limit    Count    bps     BW bps     Hysteresis    Admitted
-----   -----    -----    -----   --------   ----------    --------
1/1/0   None     1        None    None       85            0
1/1/1   None     2        15000   10000      85            2000
```

Meaning Table 14 on page 53 lists the output fields of the **show mroute port count** command.

Table 58: show mroute port count Output Fields

Field Name	Field Description
Port	Slot or port value on the router
Limit	None (reserved for future functionality)
Count	Number of multicast route outgoing interfaces on the specified port
BW bps	Bandwidth limit (in bits per second)
Priority BW bps	Priority bandwidth limit (in bits per second)
Admitted	Bandwidth admitted on the port (in bits per second)

Related Documentation

- Defining a Multicast Bandwidth Map on page 8
- Creating Mroute Port Limits on page 29
- Enabling Port-Level Admission Bandwidth Control on page 29
- show mroute port count

Monitoring IPv6 Multicast Entries in a Source or Group

Purpose Display information about the number of groups and sources.

- You can specify a multicast group address or both a multicast group address and a multicast source address to display information about a particular multicast forwarding entry.
- You can use the **active** option to display information for the active multicast routes.
- You can specify the bandwidth threshold along with the **active** option to display information for the active multicast routes with admission bandwidth greater than the specified bandwidth threshold. The default bandwidth threshold is 4000 bps.

Action To display the number of groups and sources:

```
host1#show ipv6 mroute count
IPv6 Multicast Routing Table

Counts: 2000 (S, G) entries
        0 (*, G) entries
```

Meaning Table 59 on page 222 lists the **show ipv6 mroute count** command output fields.

Table 59: show ipv6 mroute count Output Fields

Field Name	Field Description
Counts	Number of types of source group mappings: • (S,G)—Number of (S,G) entries • (*,G)—Number of (*,G) entries

Related Documentation

- show ipv6 mroute

Monitoring Join Interface Details when IPv6 OIF Mapping is Configured

Purpose Display details of the join interfaces corresponding to the mapped interfaces when OIF mapping is configured.

- You can specify a multicast group address or both a multicast group address and a multicast source address to display information about a particular multicast forwarding entry.
- You can use the **active** option to display information for the active multicast routes.

- You can specify the bandwidth threshold along with the **active** option to display information for the active multicast routes with admission bandwidth greater than the specified bandwidth threshold. The default bandwidth threshold is 4000 bps.

Action To display details of the join interfaces:

```
host1#show ipv6 mroute oif-detail
                  IPv6 Multicast Routing Table

(S, G) uptime d h:m:s[, expires d h:m:s]
[Data rate: Kbps] [SPT threshold: Kbps]
[Threshold: Kbps]
  [Admission bandwidth: bps]
  [QoS bandwidth: bps]
  RPF route: addr/mask, incoming interface
             neighbor address, owner route-owner
  Incoming interface list:
      Interface (addr/mask), State/Owner [(RPF IIF)]
  Outgoing interface list:
      Interface (addr/mask), State/Owner, Uptime/Expires
  Join interface list:
  Interface (addr/mask)[, State/Owner]

 (2000::1, ff1e::1) uptime 0 00:00:04
  Admission bandwidth: 60000000 bps
  QoS bandwidth: 0 bps (adaptive)
  RPF route: 2000::1/128, incoming interface GigabitEthernet13/1/1
             neighbor 2000::1, owner Local
  Incoming interface list:
      GigabitEthernet13/1/1 (fe80::290:1aff:fe42:8e5a/128), Accept/Mld (RPF IIF)

  Outgoing interface list:
      GigabitEthernet13/1/4 (fe80::290:1aff:fe42:8e5d/128), Blocked
(intf-adm-limit)/Mld, 0 00:00:04/never
     Join interface list:
      GigabitEthernet15/0/1.10 (fe80::290:1aff:fe42:2731/128)

Counts: 1 (S, G) entries
        0 (*, G) entries
```

Meaning Table 60 on page 223 lists the **show ipv6 mroute oif-detail** command output fields.

Table 60: show ipv6 mroute oif-detail Output Fields

Field Name	Field Description
(S,G)	IPv6 addresses of the multicast source and the multicast group
Uptime	Length of time that the (S,G) pair has been active, in *days hours:minutes:seconds* format
Expires	Length of time for which the (S,G) pair will be active, in *days hours:minutes:seconds* format
Admission bandwidth	Admission bandwidth (in bps)
QoS bandwidth	QoS bandwidth (in bps)

Table 60: show ipv6 mroute oif-detail Output Fields *(continued)*

Field Name	Field Description
RPF Route	IPv6 address and prefix of the RPF route
Incoming interface	Type and specifier of the incoming interface for the RPF route
neighbor address	IPv6 address of the neighbor
owner	Owner of the route: • Local—route belonging to the local interface • Static—Static route • Other protocols—Route established by a protocol
Incoming interface list	List of incoming interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: accept or discard • Multicast protocol that owns the interface • Time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Time that the interface will cease to be active in this multicast forwarding entry, in *days hours:minutes:seconds* format
Outgoing interface list	List of outgoing interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Forward or Blocked (intf-adm-limit, join-intf-adm-limit, port-adm-limit, port-limit, port-priority-limit) • Protocol running on the interface: PIM or MLD • Time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Time that the interface will cease to be active in this multicast forwarding entry, in *days hours:minutes:seconds* format

Table 60: show ipv6 mroute oif-detail Output Fields *(continued)*

Field Name	Field Description
Join interface list	List of join interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Forward or Blocked (intf-adm-limit, join-intf-adm-limit, port-adm-limit, port-limit, port-priority-adm-limit) • Protocol running on the interface: PIM, DVMRP, or IGMP • Amount of time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Length of time that the interface can remain active in this multicast forwarding entry, in *days hours:minutes:seconds* format or *never*
Counts	Numbers of types of source group mappings: • (S,G)—Number of (S,G) entries • (*,G)—Number of (*,G) entries

Related Documentation

- Blocking IPv6 Mroutes on page 182
- Creating IPv6 Mroute Port Limits on page 185
- show ipv6 mroute

Monitoring IPv6 Multicast Statistics

Purpose Display statistics for packets received through multicast routes that the router has added to the multicast routing table and established on the appropriate line modules.

- You can specify a multicast group address or both a multicast group address and a multicast source address to display information about a particular multicast forwarding entry.
- You can use the **active** option to display information for the active multicast routes.
- You can specify the bandwidth threshold along with the **active** option to display information for the active multicast routes with admission bandwidth greater than the specified bandwidth threshold. The default bandwidth threshold is 4000 bps.

Action To display the statistics of the multicast routes added to the multicast routing table:

```
host#show ipv6 mroute statistics
IPv6 Multicast Routing Table

(S, G) uptime d h:m:s[, expires d h:m:s]
  [Admission bandwidth: bps]
  [QoS bandwidth: bps]
  RPF route: addr/mask, incoming interface
```

```
              neighbor address, owner route-owner
  Incoming interface list:
      Interface (addr/mask), State/Owner [(RPF IIF)]
  Outgoing interface list:
      Interface (addr/mask), State/Owner, Uptime/Expires

(10:0:0:1:1::, ff0e::1) uptime 0 01:05:23
  Admission bandwidth:
  RPF route: 10:0:0:1::/64, incoming interface ATM2/3.1001
              neighbor 10:0:0:1::1, owner Local
  Incoming interface list:
      ATM2/3.1001 (10:0:0:1::1/64), Accept/Pim (RPF IIF)
  Outgoing interface list:
      ATM2/0.200 (21:2:2:21::2:1/60), Forward/Pim, 0 01:05:23/never
  Statistics:
      Received   : 346 pkts, 22144 bytes
      Forwarded  : 346 pkts, 22144 bytes
      Rcvd on OIF: 0 pkts

(10:0:0:1:2::, ff0e::1) uptime 0 01:05:23
  RPF route: 10:0:0:1::/64, incoming interface ATM2/3.1001
              neighbor 10:0:0:1::1, owner Local
  Incoming interface list:
      ATM2/3.1001 (10:0:0:1::1/64), Accept/Pim (RPF IIF)
  Outgoing interface list:
      ATM2/0.200 (21:2:2:21::2:1/60), Forward/Pim, 0 01:05:26/never
  Statistics:
      Received   : 346 pkts, 22144 bytes
      Forwarded  : 346 pkts, 22144 bytes
      Rcvd on OIF: 0 pkts
```

Meaning Table 61 on page 226 lists the **show ipv6 mroute statistics** command output fields.

Table 61: show ipv6 mroute statistics Output Fields

Field Name	Field Description
(S,G)	IPv6 addresses of the multicast source and the multicast group
Uptime	Length of time that the (S,G) pair has been active, in *days hours:minutes:seconds* format
Expires	Length of time for which the (S,G) pair will be active, in *days hours:minutes:seconds* format
Admission bandwidth	Admission bandwidth (in bps)
QoS bandwidth	QoS bandwidth (in bps)
RPF Route	IPv6 address and prefix of the RPF route
Incoming interface	Type and specifier of the incoming interface for the RPF route
neighbor address	IPv6 address of the neighbor

Table 61: show ipv6 mroute statistics Output Fields *(continued)*

Field Name	Field Description
owner	Owner of the route: • Local—route belonging to the local interface • Static—Static route • Other protocols—Route established by a protocol
Incoming interface list	List of incoming interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: accept or discard • Multicast protocol that owns the interface • Time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Time that the interface will cease to be active in this multicast forwarding entry, in *days hours:minutes:seconds* format
Outgoing interface list	List of outgoing interfaces on the router. Details include: • Type of interface and its specifier • Action that the interface takes with packets: Forward or Blocked (intf-adm-limit, join-intf-adm-limit, port-adm-limit, port-limit, port-priority-limit) • Protocol running on the interface: PIM or MLD • Time that the interface has been active in this multicast forwarding entry, in *days hours:minutes:seconds* format • Time that the interface will cease to be active in this multicast forwarding entry, in *days hours:minutes:seconds* format
Statistics	• Received—Number of packets and bytes that the virtual router received for this multicast route • Forwarded—Number of packets and statistics that the virtual router has forwarded for this multicast route • Rcvd on OIF—Number of packets and statistics that the virtual router has received on the OIF for this multicast route NOTE: The display shows statistics after the virtual router has added the multicast route to the multicast routing table and established the route on the appropriate line module. Statistics for interactions before the route is established on the line module are not displayed.

Related Documentation

- Blocking IPv6 Mroutes on page 182
- Creating IPv6 Mroute Port Limits on page 185
- show ipv6 mroute

Monitoring Summary Information of IPv6 Multicast Routes

Purpose

Display a summary of all or specified multicast routes.

- You can specify a multicast group address or both a multicast group address and a multicast source address to display information about a particular multicast forwarding entry.
- You can use the **active** option to display information for the active multicast routes.
- You can specify the bandwidth threshold along with the **active** option to display information for the active multicast routes with admission bandwidth greater than the specified bandwidth threshold. The default bandwidth threshold is 4000 bps.

Action

To display a summary of all multicast routes:

```
host1#show ipv6 mroute summary
IPv6 Multicast Routing Table

Group Address   Source Address  RPF route          RPF Iif         #Oifs
--------------- --------------- ------------------ --------------- -----
ff0e::1         10:0:0:1:1::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:2::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:3::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:4::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:5::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:6::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:7::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:8::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:9::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:a::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:b::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:c::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:d::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:e::    10:0:0:1::/64      ATM2/3.1001     1
ff0e::1         10:0:0:1:f::    10:0:0:1::/64      ATM2/3.1001     1

Counts: 16 (S, G) entries
        0 (*, G) entries
```

Meaning

Table 62 on page 228 lists the **show ipv6 mroute summary** command output fields.

Table 62: show ipv6 mroute summary Output Fields

Field Name	Field Description
Group Address	IP address of the multicast group
Source Address	IP address of the multicast source
RPF Route	IP address and network mask of the RPF route

Table 62: show ipv6 mroute summary Output Fields *(continued)*

Field Name	Field Description
RPF Iif	Type and identifier for the incoming interface for the RPF route
#Oifs	Number of outgoing interfaces
Counts	Numbers of types of (S,G) mappings: • (S,G)—Number of (S,G) entries • (*,G)—Number of (*,G) entries

Related Documentation

- Defining IPv6 Static Routes for Reverse-Path Forwarding on page 165
- show ipv6 mroute

Monitoring IPv6 Multicast Protocols Enabled on the Router

Purpose Display information about multicast protocols enabled on the router.

Action To display information about enabled multicast protocols:

```
host1:2#show ipv6 multicast protocols
Multicast protocols:

Protocol Pim
    Type: Sparse
    Interfaces: 1 registered, 1 owned
    Registered interfaces:
        ATM2/1.103 (21:2:2:22::1:2/60) owner Pim
Protocol Mld
    Type: Local
    Interfaces: 1000 registered, 1000 owned
    Registered interfaces:
        ATM2/0.131 (31:2:2:22::2:2/604) local Mld owner Mld
            Admission-bandwidth 2000000/10000000 bps
            QoS Adjust 2000 bps
        ATM2/0.132 (31:2:2:22::2:3/60) local Mld owner Mld
            Admission-bandwidth 0/10000000 bps
            QoS Adjust 0 bps
        ATM2/0.133 (31:2:2:22::2:4/60) local Mld owner Mld
            Admission-bandwidth 8000000/10000000 bps, 2 Blocked OIFs
            QoS Adjust 0 bps
...
Count: 2  protocols
```

Meaning Table 63 on page 229 lists the **show ipv6 multicast protocols** command output fields.

Table 63: show ipv6 multicast protocols Output Fields

Field Name	Field Description
Protocol	Name of the multicast protocol

Table 63: show ipv6 multicast protocols Output Fields *(continued)*

Field Name	Field Description
Type	Mode of the multicast protocol: • For PIM—Sparse • For MLD—Local
Interfaces	• registered—Number of interfaces on which the protocol is configured • owned—Number of interfaces that a protocol owns. If you configure only MLD on an interface, MLD owns the interface. However, if you configure MLD and PIM on the same interface, PIM owns the interface.
Registered interfaces	Includes the following information about interfaces on which the protocol is configured: • Types and identifiers of interfaces. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide* . • Protocols configured on the interface and the protocol that owns the interface. If you configure only MLD on an interface, MLD owns the interface. However, if you configure MLD and PIM on the same interface, PIM owns the interface. • Admitted bandwidth / configured admission bandwidth (in bps) • Number of blocked OIFs • QoS adjustment bandwidth (in bps)
Count	Number of multicast protocols on the virtual router

Related Documentation

- Example: Configuring an IPv6 Multicast Bandwidth Map on page 171
- Activating IPv6 Multicast QoS Adjustment Functions on page 175
- Enabling Interface-Level Admission Bandwidth Limitation for IPv6 on page 183
- Enabling Port-Level Admission Bandwidth Limitation for IPv6 on page 185
- show ipv6 multicast protocols

Monitoring Summary Information of IPv6 Multicast Protocols Enabled on the Router

Purpose Display information about multicast protocols enabled on the router.

Action To display a summary of information about multicast protocols enabled on the router:

```
host1#show ipv6 multicast protocols brief
Protocol   Registered Owned       Type
           Interfaces Interfaces
```



```
--------- ---------- ---------- -------------------
Pim              1          1    Sparse
Mld              1          1    Local

Count: 2 protocols
```

Meaning Table 64 on page 231 lists the **show ipv6 multicast protocols brief** command output fields.

Table 64: show ipv6 multicast protocols brief Output Fields

Field Name	Field Description
Protocol	Name of the multicast protocol
Registered Interfaces	Number of interfaces on which the protocol is configured
Owned Interfaces	Number of interfaces that a protocol owns. If you configure only MLD on an interface, MLD owns the interface. However, if you configure MLD and PIM on the same interface, PIM owns the interface
Type	Mode of the multicast protocol: • For PIM—Sparse • For MLD—Local
Count	Number of multicast protocols on the virtual router

Related Documentation

- show ipv6 multicast protocols

Monitoring IPv6 Multicast Status on a Virtual Router

Purpose Display information about the status of IPv6 multicast on the virtual router.

Action To display information about the status of IPv6 multicast on the virtual router:

```
host1#show ipv6 multicast routing
Multicast forwarding is enabled on this router
Multicast graceful restart is complete (timer 0 seconds) on this router
Multicast cache-miss processing is enabled on this router


Multicast forwarding is enabled on this router
Multicast graceful restart is complete (timer 0 seconds) on this router
Multicast cache-miss processing is enabled on this router
```

Related Documentation

- Enabling IPv6 Multicast on page 165
- show ipv6 multicast routing

CHAPTER 11

Configuring MLD and MLD Proxy

Hosts use Multicast Listener Discovery (MLD) protocol in IPv6 to report their multicast group memberships to neighboring routers. Similarly, multicast routers, such as the E Series router, use MLD to discover which of their hosts belong to multicast groups.

This chapter describes how to configure MLD and MLD proxy on an E Series router; it contains the following sections:

MLD Overview

The IPv6 address scheme uses hexadecimal FF at the start of an address for IPv6 multicast. Multicast Listener Discovery (MLD) is a protocol that uses these addresses. The following addresses have specific functions:

- You can assign only multicast addresses of global-scope (that is, containing an FF*x*E prefix, where *x* is the flags field) to a multicast group.
- FF02::1 is the link-scope all-nodes address—A packet sent to this address reaches all nodes on a subnetwork.

- FF02::2 is the link-scope all-routers address—A packet sent to this address reaches all routers on a subnetwork.
- FF02::16 is the link-scope all-MLDv2 routers address—A packet sent to this address reaches all MLDv2 routers on a subnetwork.

This implementation of MLD complies with MLD versions 1 and 2. MLDv2 allows for source-specific join and leave messages and is backward compatible with MLDv1. Configuring MLDv1 with the SSM mapping feature provides support for source-specific joins.

MLDv1 mode interfaces exchange the following types of messages between routers and hosts:

- Multicast Listener Queries on page 234
- Multicast Listener Reports on page 235
- Multicast Listener Done Messages on page 235

MLDv2 mode interfaces exchange the following types of messages with MLDv2 hosts:

- Multicast Listener Queries on page 234
- MLDv2 “Multicast Listener Reports” on page 235

Multicast Listener Queries

A multicast router can be a querier or a nonquerier. There is only one querier on a network at any time. Multicast routers monitor queries from other multicast routers to determine the status of the querier. If the querier hears a query from a router with a lower IPv6 address, it relinquishes its role to that router.

MLDv1 and MLDv2 mode interfaces send two types of multicast listener queries to hosts on the network:

- General queries to the all-nodes address (FF02::1)
- Specific queries to the appropriate multicast group address

MLDv2 mode interfaces send the following type of queries to MLDv2 hosts:

- General queries
- Group-specific queries
- Source-specific queries

The purpose of a membership group query is to discover the multicast groups to which a host belongs.

MLDv1 and MLDv2 multicast listener queries have a Max Response Time field. This response time is the maximum that a host can take to reply to a query.

Multicast Listener Reports

When a host receives a multicast listener query, it identifies the groups associated with the query and determines to which groups it belongs. The host then sets a timer, with a value less than the Max Response Time field in the query, for each group to which it belongs.

When the timer expires, the host sends a multicast listener report to the group address. When a multicast router receives a report, it adds the group to the membership list for the network and sets a timer to the *multicast address listening interval*. If this timer expires before the router receives another multicast listener report, the router determines that the group has no members left on the network.

If the router does not receive any reports for a specific multicast group within the *maximum response time*, it determines that the group has no members on the network. The router does not forward subsequent multicasts for that group to the network.

MLDv2 supports an extended report format that allows you to report multiple groups and source lists in a single report. These reports are addressed to the all-MLDv2 router's multicast address (FF02::16).

Multicast Listener Done Messages

When an MLDv1 host leaves a group, it sends a multicast listener done message to multicast routers on the network. A host generally addresses multicast listener done messages to the all-routers address, FF02::2.

When an MLDv2 host leaves a group, it sends a multicast listener report. This report includes an empty source list for that group.

Related Documentation

- MLD Platform Considerations on page 235
- MLD References on page 236

MLD Platform Considerations

For information about modules that support MLD on the ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support MLD.

For information about modules that support MLD on the E120 and E320 Broadband Services Routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support MLD.

Related Documentation

- MLD Overview on page 233
- MLD References on page 236

MLD References

For more information about MLD, see the following resources:

- RFC 3710—Multicast Listener Discovery (MLD) for IPv6 (October 1999) on page 578
- IGMP/MLD-based Multicast Forwarding ('IGMP/MLD Proxying')—draft-ietf-magma-igmp-proxy-06.txt (October 2004 expiration) on page 587
- Multicast Group Membership Discovery MIB—draft-ietf-magma-mgmd-mib-06.txt (October 2004 expiration) on page 587

Related Documentation

- MLD Overview on page 233
- MLD Platform Considerations on page 235

Static and Dynamic MLD Interfaces

The router supports *static* and *dynamic* Multicast Listener Discovery (MLD) interfaces. Unlike static interfaces, dynamic interfaces are not restored when you reboot the router. For some protocols, dynamic layers can build on static layers in an interface; however, in a dynamic MLD interface, all the layers are dynamic. See Figure 19 on page 236 for examples of static and dynamic MLD interfaces.

Figure 19: Static and Dynamic MLD Interfaces

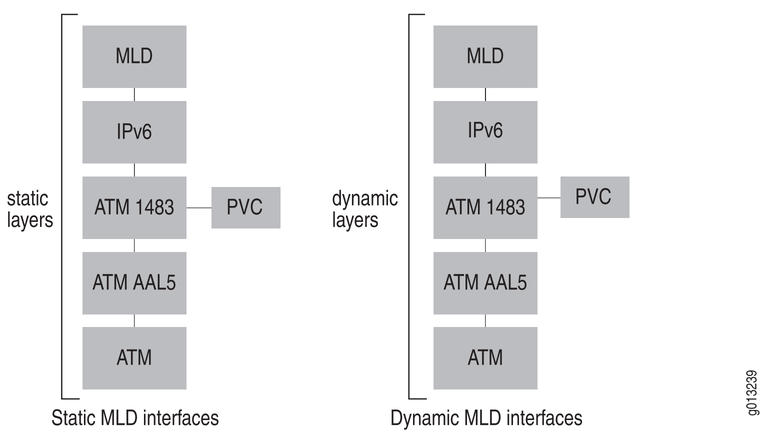

Static MLD interfaces are configured with software such as the command-line interface (CLI) or an SNMP application; dynamic MLD interfaces are configured with a profile. A profile comprises a set of attributes for an interface; a profile for dynamic MLD interfaces contains attributes for configuring all the layers in the interface.

You define a profile by using the same CLI commands that you use to configure a static MLD interface; however, the mode in which you use the commands differs. Use the commands in Interface Configuration mode to configure a static MLD interface and in Profile Configuration mode to define a profile.

When you have defined a profile, you can apply it to an interface or a group of interfaces. Profiles provide an efficient method of creating and managing large numbers of dynamic interfaces. For detailed information about creating and assigning profiles, see *Configuring Dynamic Interfaces* in the *JunosE Link Layer Configuration Guide*. When you create a profile for dynamic MLD interfaces, specify attributes for configuring all layers in the interface.

You use the MLD commands shown in Table 65 on page 237 to configure a static MLD interface. You also use these commands to define the attributes for the MLD layer when you create a profile for dynamic MLD interfaces.

Table 65: Static MLD Commands

ipv6 mld	ipv6 mld query-interval
ipv6 mld access-group	ipv6 mld query-max-response-time
ipv6 mld access-source-group	ipv6 mld robustness
ipv6 mld explicit-tracking	ipv6 mld static-include
ipv6 mld group limit	ipv6 mld static-exclude
ipv6 mld immediate-leave	ipv6 mld static-group
ipv6 mld last-member-query-interval	ipv6 mld version
ipv6 mld querier-timeout	

You can also use various MLD-specific RADIUS attributes in RADIUS Access-Accept messages as an alternative method of configuring certain values. See Juniper Networks VSAs Supported for Subscriber AAA Access Messages for additional information.

Related Documentation

- MLD Overview on page 233
- MLD Platform Considerations on page 235
- MLD References on page 236

Enabling MLD on an Interface

You must start MLD on each interface that you want to use the protocol. You can configure MLD and PIM on the same interface. If you configure only MLD on an interface, the router determines that MLD owns that interface. If you configure MLD and PIM on an interface, the router determines that PIM owns the interface.

In an MLDv1 or MLDv2 network, the querier is the router with the lowest IPv6 address.

NOTE: You can configure MLD only on IPv6 interfaces. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

To start MLD and set the MLD version on the interface, complete the following steps:

1. Enable MLD on the interface (MLDv2 is the default version).

   host1:boston(config-if)#**ipv6 mld**

   Use the **no** version to disable MLD on an interface.
2. (MLDv1) Specify the MLD version for the interface.

   host1:boston(config-if)#**ipv6 mld version 1**

   Use the **no** version to set the version to the default, MLDv2.

Related Documentation

- Monitoring MLD Information on a Virtual Router on page 255
- Monitoring MLD Groups on page 256
- Monitoring MLD Interfaces on page 258

- ipv6 mld
- ipv6 mld version

Configuring MLD Settings on an Interface

When you start MLD on an interface, it operates with the default settings. You can, however, modify:

- The method that the router uses to remove hosts from multicast groups
- The time interval at which the querier sends multicast listener queries
- The time that a querier waits before sending a new query to hosts from which it receives multicast listener done messages
- The time that a non-querier waits for queries from the current querier before sending query messages to assume responsibility of querier
- The time that a host can take to reply to a query (maximum response time)
- The number of times that the router sends each MLD message from this interface

NOTE: You can configure MLD only on IPv6 interfaces. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

To configure MLD settings for an interface:

- Enable the immediate-leave feature.

  host1:boston(config-if)#**ipv6 mld immediate-leave**

  The router immediately removes a host from a multicast group after receiving a multicast listener done message from the host associated with this interface.

  Use the **no** version to restore the default behavior, in which the router removes a host from a multicast group if that host does not return a multicast listener report within a certain length of time after receiving a multicast listener query from the router.

  CAUTION: Issue this command only on MLD interfaces to which one MLD host is connected. If more than one MLD host is connected to a LAN through the same interface, and one host sends a done message, the router removes all hosts on the interface from the multicast group. The router loses contact with the hosts that should remain in the multicast group until they send join requests in response to the router's next general multicast listener query.

  NOTE: As an alternative method, you can configure the immediate-leave feature using the MLD-Immediate-Leave RADIUS attribute (VSA 26-100) in RADIUS Access-Accept messages. The RADIUS setting takes precedence over a CLI setting. The **radius ignore** command is used to ignore or accept RADIUS attributes from Access-Accept messages.

- Configure the last member query interval, which specifies the maximum time the router will wait for a response after sending a last member query. Using a lower value allows members to leave groups more quickly.

  host1:boston(config-if)#**ipv6 mld last-member-query-interval 90**

  Use the **no** version to restore the default, 10-tenths of a second (1 second).

- Configure the time that the interface waits before declaring itself as the querier.

  host1:boston(config-if)#**ipv6 mld querier-timeout 200**

  Use the **no** version to set the time to the default, twice the query interval.

- Configure the query interval, which specifies how often the router sends MLD host-query packets from the interface.

  host1:boston(config-if)#**ipv6 mld query-interval 100**

  Use the **no** version to set the polling interval to the default, 125 seconds.

  NOTE: As an alternative method, you can configure the query interval using the MLD-Query-Interval RADIUS attribute (VSA 26-98) in RADIUS Access-Accept messages. The RADIUS setting takes precedence over a CLI setting. The **radius ignore** command is used to ignore or accept RADIUS attributes from Access-Accept messages.

- Configure the maximum response time, which specifies the period during which the host is expected to respond to an MLD query. MLDv1 and MLDv2 include this value in MLD query messages sent out on the interface. Using a lower value allows members to join and leave groups more quickly.

  host1:boston(config-if)#ipv6 mld query-max-response-time 120

  Use the **no** version to restore the default, 100 tenths of a second (10 seconds).

  NOTE: As an alternative method, you can configure the query maximum response time using the MLD-Query-Max-Resp-Time RADIUS attribute (VSA 26-99) in RADIUS Access-Accept messages. The RADIUS setting takes precedence over a CLI setting. The **radius ignore** command is used to ignore or accept RADIUS attributes from Access-Accept messages.

- Configure the robustness value, which specifies the number of times that the router sends MLD group-specific queries before declaring a group to no longer have any members on the interface.

  host1:boston(config-if)#ipv6 mld robustness 2

  Use the **no** version to restore the default, 3.

Related Documentation

- Monitoring MLD Interfaces on page 258
- Monitoring Summary Information for MLD Interfaces on page 262
- Juniper Networks VSAs Supported for Subscriber AAA Access Messages
- CLI Commands Used to Include or Exclude Attributes in RADIUS Messages
- ipv6 mld immediate-leave
- ipv6 mld last-member-query-interval
- ipv6 mld querier-timeout
- ipv6 mld query-interval
- ipv6 mld query-max-response-time
- ipv6 mld robustness
- radius ignore

Configuring Multicast Groups for MLD

You can configure various attributes related to MLD multicast groups. This topic describes how to configure multicast group attributes.

NOTE: You can configure MLD only on IPv6 interfaces. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Specifying MLD Multicast Groups

You can use a standard IPv6 access list to specify the multicast groups that a host can join.

To restrict the host to join only the groups permitted by the IPv6 access list:

- Restrict hosts on this subnetwork to join only multicast groups on the specified IPv6 access list. The access list is queried whenever the router receives an MLDv1 report requesting membership of a group and MLDv2 ChangeToInclude, IsInclude, ChangeToExclude, or IsExclude reports. The request is ignored if the access list query fails. The **ipv6 mld access-group** command uses IPv6 access lists, which allow both source and destination/group addresses to be specified. You must set the source address to any.

  host1:boston(config-if)#**ipv6 mld access-group boston-list**

  Use the **no** version to dissociate the interface from an access list and to allow hosts on the interface to join any multicast group.

  NOTE: As an alternative method, you can restrict the hosts using the MLD-Access-Name RADIUS attribute (VSA 26-74) in RADIUS Access-Accept messages. The **radius ignore** command is used to ignore or accept RADIUS attributes from Access-Accept messages.

- Restrict hosts on this subnetwork to membership only in (S,G) pairs (also known as channels) permitted by the specified IPv6 access list. When configured, both source and group addresses query the associated access list whenever the router receives an MLDv2 report requesting membership of the (S,G) pairs (that is, the router receives an MLDv2 ChangeToInclude, IsInclude, or AllowNewSource group report). The request is ignored if the access list query fails. The **ipv6 mld access-source-group** command uses IPv6 access lists, which allow both source and destination group addresses to be specified.

  host1:boston(config-if)#**ipv6 mld access-source-group dallas-list**

  Use the **no** version to remove any access list restriction.

NOTE: As an alternative method, you can restrict the hosts using the MLD-Access-Src-Name RADIUS attribute (VSA 26-75) in RADIUS Access-Accept messages. The **radius ignore** command is used to ignore or accept RADIUS attributes from Access-Accept messages.

Assigning an MLD Multicast Group to an Interface

You can assign an interface to send and receive all traffic for a particular multicast group. This feature allows you to control MLD traffic and to test the behavior of multicast protocols in the network.

To send and receive all traffic for a multicast group from the interface:

- Issue the **ipv6 mld static-group** command in Interface Configuration mode. The interface sets no timers for this group.

  ```
  host1:boston(config-if)#ipv6 mld static-group ff0e::1
  ```

  Use the **no** version to remove the group from the interface.

Configuring MLD Group Outgoing Interface Mapping

You can configure an MLD protocol interface to use a different OIF for multicast-data-forwarding by applying an OIF map. When you configure an OIF map on an MLD protocol interface, the map is applied to all MLD membership requests that the interface receives. To configure OIF mapping on an interface, you first create the OIF map using the **ipv6 mld oif-map** command and then apply that map to an interface with the **ipv6 mld apply-oif-map** command.

To properly configure an interface used in the OIF map for multicast-data-forwarding capability, you must configure the interface version as passive with the **ipv6 mld version** command. You can either specify a passive interface as the OIF or specify the OIF as *self* (to use the MLD protocol interface as the OIF) in the **ipv6 mld oif-map** command.

To configure group outgoing interface mapping:

1. Create an OIF map.

   ```
   host1(config)#ipv6 mld oif-map OIFMAP atm 3/0.1 ff0e::1:1/128 2001::1:1/128
   host1(config)#ipv6 mld oif-map OIFMAP atm 3/0.2 ff0e::1:1/128 2001::1:1/128
   host1(config)#ipv6 mld oif-map OIFMAP atm 3/0.3 ff0e::1:1/128
   host1(config)#ipv6 mld oif-map OIFMAP atm 3/0.4 ff0e::1:0/112 2001::1:1/128
   host1(config)#ipv6 mld oif-map OIFMAP atm 3/0.5 ff0e::1:0/112 2001::1:1/128
   host1(config)#ipv6 mld oif-map OIFMAP self ::/0 2001::1:0/112
   ```

   Use the **no** version to remove an outgoing interface map attribute.

NOTE: As an alternative method, you can create the OIF map using the MLD-OIF-Map-Name RADIUS attribute (VSA 26-76) in RADIUS Access-Accept messages. The **radius ignore** command is used to ignore or accept RADIUS attributes from Access-Accept messages.

2. Apply the created OIF map to the current interface.

   host1(config-subif)#**ipv6 mld apply-oif-map OIFMAP**

   Use the **no** version to remove the outgoing interface map association from the interface.

3. Configure the MLD version for the interface as passive with only multicast-data-forwarding capability.

   host1:dallas(config-if)#**ipv6 mld version passive**

   Use the **no** version to set the version to the default, MLDv2.

NOTE: As an alternative method, you can configure the MLD version using the MLD-Version RADIUS attribute (VSA 26-77) in RADIUS Access-Accept messages. The **radius ignore** command is used to ignore or accept RADIUS attributes from Access-Accept messages.

Related Documentation

- Overview of Limiting the Number of Accepted MLD Groups on page 244
- Limiting the Number of Accepted MLD Groups on page 247
- Monitoring MLD Groups on page 256

- Juniper Networks VSAs Supported for Subscriber AAA Access Messages
- CLI Commands Used to Include or Exclude Attributes in RADIUS Messages
- ipv6 mld access-group
- ipv6 mld access-source-group
- ipv6 mld apply-oif-map
- ipv6 mld oif-map
- ipv6 mld static-group
- ipv6 mld version
- radius ignore

MLD SSM Mapping Overview

Source-specific multicast (SSM) mapping enables the router to determine one or more source addresses for group G. The mapping effectively translates a Multicast Listener Discovery version 1 (MLDv1) multicast listener report to an MLDv2 report, enabling the router to continue as if it had initially received an MLDv2 report. After the router is joined to these groups, it sends out Protocol Independent Multicast (PIM) join messages and continues to enable joining from these groups, as long as it continues to receive MLDv1 membership reports and no change occurs to the SSM mapping for the group.

When you statically configure SSM mapping, the router can discover source addresses from a statically configured table.

The following applies when you configure SSM mapping:

- When an SSM map is configured without any matching access list, SSM mapping is not applied on the incoming (*,G) groups. The PIM SSM range must deny any unacceptable SSM group addresses.

NOTE: **An access list must be explicitly configured with the same name as that of the SSM map and group addresses that are to be SSM mapped.**

- When you issue the **no ipv6 mld ssm-map enable** command, the router removes all SSM map (S,G) states and establishes a (*,G) state.
- You can enter multiple **ipv6 mld ssm-map static** commands for different access lists. Also, you can enter multiple **ipv6 mld ssm-map static** commands for the same access list, as long as the access list uses different source addresses.
- SSM maps do not process statically configured groups.

Related Documentation

- Configuring MLD SSM Mapping on page 247
- ipv6 mld ssm-map enable
- ipv6 mld ssm-map static

Overview of Limiting the Number of Accepted MLD Groups

By default, there is no limit on the number of MLD groups that an MLD interface can accept. However, you can manage multicast traffic on the router by restricting the number of MLD groups accepted by:

- A specific port on an I/O module
- A specific MLD interface

If you set limits for both a port and interfaces on that port, the router uses the lower of the two limits when determining how many MLD groups an interface can accept. For example, if you set a limit of 10 groups for the port and 15 groups for each interface, the router allows only 10 groups to be accepted among the interfaces.

However, if you set a limit for a port and that limit is lower than the number of groups currently accepted by the interfaces on that port, the router does not dissociate the groups from the interfaces. The router enforces the new limit on the port when the number of groups associated with the interfaces falls to that limit. For example, if the interfaces on the port have accepted a total of 15 groups, and you set a limit of 10 groups on the port, the router does not disconnect any of the groups and does not allow the interfaces to accept any more groups. Over time, some groups leave the interfaces and, eventually, a maximum of ten groups remains connected.

Related Documentation

- Configuring Multicast Groups for MLD on page 240
- Limiting the Number of Accepted MLD Groups on page 247

MLD Traffic Overview

MLDv2 extends MLDv1 functionality with the ability to include or exclude specific multicast traffic sources. That is, with MLDv2, hosts signal (S,G) pairs that they want to include or exclude.

For hosts that cannot signal group membership dynamically, you can use the **ipv6 mld static-include** or **ipv6 mld static-exclude** command to statically include or exclude multicast traffic, respectively.

MLDv2 is the industry-designated standard protocol for hosts to signal channel subscriptions in SSM. For additional information about SSM, see *PIM Source-Specific Multicast* in “Understanding PIM for IPv4 Multicast” on page 90.

Related Documentation

- Including and Excluding MLD Traffic on page 248

MLD Explicit Host Tracking Overview

Explicit host tracking enables the router to explicitly track each individual host that is joined to a group or channel on a particular multi-access network.

Explicit host tracking provides the following benefits:

- Minimal leave latency when a host leaves a multicast group or channel. When the router receives a leave message for a group or channel on an interface, it accesses a list of hosts and immediately stops forwarding traffic if the sender is the last host to request traffic for that group or channel. The leave latency is bound only by the packet transmission latencies in the multi-access network and the processing time in the router.
- Ability to change channels quickly in networks where bandwidth is constrained between a multicast-enabled router and hosts.
- Ability to determine what multicast hosts are joined to particular multicast groups or channels; this is useful for accounting purposes.
- Reduction of control message traffic on the network because, when it receives a leave message, the router no longer needs to send out MLD queries to verify membership. As a result, interested hosts also do not need to respond to these queries with reports.
- Tracking based on MLD reports for hosts in both include and exclude modes for every multicast group or channel on an interface.

When the router is configured for explicit host tracking and starts performing immediate leave using the host information collected, every leave message received for a group or channel is treated as follows:

- The router checks the number of hosts that receive traffic from this group or channel.
- If the host sending the leave message is the only host, it performs immediate leave for that group or channel on that interface. The router removes the interface from the multicast group or channel immediately, without sending out a group or group-source specific query and waiting for the last member query interval.
- If the host sending the leave message is not the only host receiving traffic for that group or channel, the router removes the host from the list of hosts on that interface, but keeps the interface in the outgoing interface list for the multicast group or channel. No group or group-source specific queries are sent.

You can enable Multicast Listener Discovery (MLD) explicit host tracking on an interface only if MLD V1 or V2 has been previously enabled on the interface. Explicit host tracking is not enabled by default when you enable MLD on the interface. Explicit host tracking cannot be configured on passive MLD interfaces.

When you enable explicit host tracking on an interface that has a membership state, the router does not immediately start performing immediate leave. For a maximum of group membership interval seconds, the router only performs host tracking. Any leave messages that the router receives during this period receive normal leave processing. Any leave messages received after this interval has elapsed receive immediate leave processing, when appropriate.

When explicit host tracking is enabled on an MLD V2 interface, even if a group has to downgrade to MLD V1 due to the presence of an MLD V1 host, explicit host tracking continues for that group. To avoid this, you can use the **disable-if-mld-v1-detected** keyword. If you select this option, the router turns off explicit host tracking for the group when MLD V1 host reports are received for the group on that interface. This option does not have any significance on an interface configured for MLD V1 and is ignored if selected.

If you execute the command on an interface that was previously enabled for immediate-leave, the configuration is accepted, immediate-leave is turned off and an appropriate warning message logged. Any attempt to configure immediate-leave on an interface that has explicit host tracking enabled is rejected and an error message logged.

Related Documentation

- Configuring MLD Explicit Host Tracking on page 248

Configuring MLD Attributes

You can configure MLD Attributes using the following set of tasks:

NOTE: You can configure MLD only on IPv6 interfaces. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Configuring MLD SSM Mapping

SSM mapping statically assigns sources to MLDv1 groups. You must use SSM mapping for MLDv1 hosts to interoperate with PIM SSM. SSM mapping allows the router to use a statically configured list to translate (*,G) memberships to (S,G) memberships.

The **ipv6 mld ssm-map static** command uses IPv6 access lists, which allow both source and destination/group addresses to be specified. You must set the source address to "any".

To statically configure SSM mapping:

1. Enable SSM mapping on the router.

   host1:boston(config)#**ipv6 mld ssm-map enable**

   Use the **no** version to disable the SSM map.

2. Specify an access list and source address for use in SSM mapping.

   host1:boston(config)#**ipv6 mld ssm-map static boston-list 2001::1**

   Use the **no** version to remove the SSM map association.

   NOTE: To operate correctly, the static source addresses must fall within the configured PIM SSM range.

Limiting the Number of Accepted MLD Groups

You can manage multicast traffic on the router by restricting the number of MLD groups.

To limit the number of accepted MLD groups:

1. Configure the maximum number of MLD groups that an interface can accept.

   host1:boston(config-if)#**ipv6 mld group limit 5**

   Use the **no** version to restore the default situation, in which there is no limit on the number of MLD groups that an interface can accept.

2. Configure the maximum number of MLD groups that a port can accept.

- On ERX models:

  host1(config)#**multicast group port 3/0 limit 5**

- On E120 and E320 routers:

  host1(config)#**multicast group port 3/1/0 limit 5**

NOTE: You can specify the identifier for the port in *slot/port* format (ERX routers) or in *slot/adapter/port* format (E120 and E320 routers) and the maximum number of MLD groups that interfaces can accept.

- slot—Number of the chassis slot in the range 0–6 (ERX7xx models), 0–13 (ERX14xx models), 0–5 (E120 router), or 0–16 (E320 router)
- adapter—Number of the bay in which the I/O adapter (IOA) resides. This identifier applies to the E120 and E320 routers only. In the software, adapter 0 identifies the right IOA bay (E120 router) and the upper IOA bay (E320 router); adapter 1 identifies the left IOA bay (E120 router) and the lower IOA bay (E320 router).
- port—Port number on the I/O module or IOA

Use the **no** version to restore the default situation, in which there is no limit on the number of MLD groups that a port can accept.

Including and Excluding MLD Traffic

MLD allows you to include or exclude specific multicast traffic sources.

To include and exclude traffic:

- Configure the interface to statically include the MLD (S,G) membership for a host that is not capable of dynamically signaling group membership.

  host1:boston(config-if)#**ipv6 mld static-include 2001::1 ff0e::1**

  Use the **no** version to remove the static designation.

- Configure the interface to statically exclude the MLD (S,G) membership for a host that is not capable of dynamically signaling group membership.

  host1:boston(config-if)#**ipv6 mld static-exclude 2001::1 ff0e::1**

  Use the **no** version to remove the static designation.

Configuring MLD Explicit Host Tracking

Explicit host tracking enables the router to explicitly track each individual host that is joined to a group or channel on a particular multi-access network.

To enable explicit host tracking on the interface:

1. Configure the MLD version for the interface.

```
host1(config)#interface 3/0.101
host1(config-if)#ipv6 mld version 2
```

2. Enable explicit host tracking on the interface without the capability of being automatically disabled even if MLD V1 hosts are detected on the MLD V2 interface. This tracking continues for a group even if the group downgrades to MLD V1 due to the presence of an MLD V1 host.

   ```
   host1(config)#ipv6 mld explicit-tracking
   ```

   Use the **no** version to disable explicit host tracking on the interface.

3. Enable explicit host tracking on the interface with the capability of being automatically disabled if MLD V1 hosts are detected on the MLD V2 interface. The router turns off explicit host tracking for the group when MLD V1 host reports are received for the group on the MLD V2 interface.

   ```
   host1(config)#ipv6 mld explicit-tracking disable-if-mld-v1-detected
   ```

   Use the **no** version with the **disable-if-mld-detected** keyword to revert to the default explicit host tracking behavior.

Disabling and Removing MLD

You can disable and reenable MLD on the virtual router. You can also remove MLD from the virtual router and re-create it on the virtual router.

To disable and enable MLD on a virtual router:

1. Disable MLD on a virtual router.

   ```
   host1(config)#virtual-router boston
   host1:boston(config)#router mld
   host1:boston(config-router)#mld disable
   ```

   Use the **no** version to enable MLD on a virtual router.

2. Create and enable MLD on a virtual router.

   - Using the **router mld** command:

     ```
     host1(config)#virtual-router boston
     host1:boston(config)#router mld
     ```

     Use the **no** version to delete MLD and MLD proxy from the virtual router.

   - Using the **ipv6 router mld** command:

     ```
     host1(config)#virtual-router boston
     host1:boston(config)#ipv6 router mld
     ```

     Use the **no** version to delete MLD and MLD proxy from the virtual router.

Related Documentation

- Configuring Multicast Groups for MLD on page 240
- MLD SSM Mapping Overview on page 243
- Overview of Limiting the Number of Accepted MLD Groups on page 244
- MLD Traffic Overview on page 245

- MLD Explicit Host Tracking Overview on page 245
- interface
- ipv6 mld explicit-tracking
- ipv6 mld group limit
- ipv6 mld ssm-map enable
- ipv6 mld ssm-map static
- ipv6 mld static-exclude
- ipv6 mld static-include
- ipv6 mld version
- ipv6 router mld
- mld disable
- multicast group port limit
- router mld
- virtual-router

MLD Proxy Overview

Multicast Listener Discovery (MLD) proxy enables the router to issue MLD host messages on behalf of hosts that the router discovered through standard MLD interfaces. The router acts as a *proxy* for its hosts. The E Series router supports MLD proxy versions 1 and 2.

Figure 20 on page 251 shows a router in an MLD proxy configuration. You enable MLD proxy on one interface, which connects to a router closer to the root of the tree. This interface is the *upstream interface*. The attached upstream router on the upstream interface should be running MLD.

You enable MLD on the interfaces that connect the router to its hosts that are farther away from the root of the tree. These interfaces are known as *downstream interfaces*.

Figure 20: Upstream and Downstream Interfaces

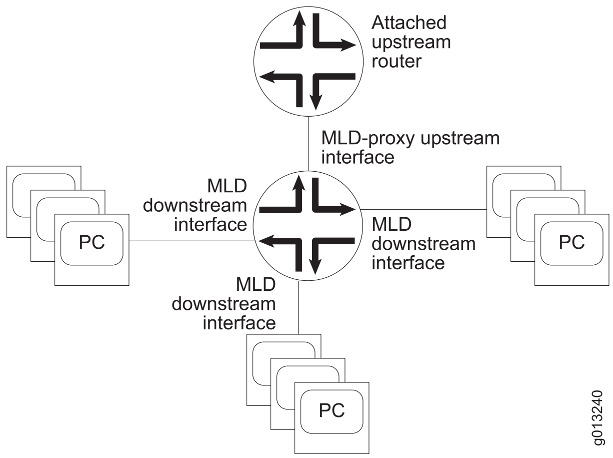

As described in “MLD Overview” on page 233, hosts interact with the router through the exchange of MLD messages. Similarly, when you configure MLD proxy, the router interacts with the router on its upstream interface through the exchange of MLD messages. However, when acting as the proxy, the router performs the host portion of the MLD task on the upstream interface as follows:

- When queried, sends multicast listener reports to the group.
- When one of its hosts joins a multicast address group to which none of its other hosts belong, sends unsolicited multicast listener reports to that group.
- When the last of its hosts in a particular multicast group leaves, the group sends either an unsolicited multicast listener done report to the all-routers address (FF02::2) for MLDv1 or an MLDv2 multicast listener report to the all-MLDv2 routers address (FF02::16).

Related Documentation

- MLD Overview on page 233
- Configuring MLD Proxy on page 251
- Establishing the MLD Proxy Baseline on page 252

Configuring MLD Proxy

To configure a downstream interface, enable MLD on that interface. To configure MLD proxy on the router, complete the following tasks:

1. Enable IPv6 multicast.
2. Identify the interface that you want to act as the upstream interface.
3. Enable MLD proxy on that interface.
4. (Optional) Specify how often the router should send unsolicited reports to routers on the upstream interface.

NOTE: You can configure MLD proxy only on IPv6 interfaces. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

To configure MLD proxy on an interface:

1. Enable MLD proxy on the upstream interface.

   host1(config-if)#**ipv6 mld-proxy**

   Use the **no** version to disable MLD proxy on an interface.

   NOTE: You can enable only one upstream interface.

2. (Optional) Configure unsolicited report interval, which specifies how often the upstream interface should transmit unsolicited reports.

   host1(config-if)#**ipv6 mld-proxy unsolicited-report-interval 600**

   Use the **no** version to transmit unsolicited reports using the default value, 100-tenths of a second (10 seconds).

   NOTE: Issue this command only on the upstream interface. Otherwise, this command has no effect.

3. (Optional) Configure the MLD proxy version for the interface.

   host1(config-if)#**ipv6 mld-proxy version 1**

   Use the **no** version to set the version to its default value, MLDv2.

Related Documentation

- MLD Proxy Overview on page 250
- Monitoring MLD Proxy Parameters on page 268
- Monitoring MLD Proxy Interfaces on page 270

- ipv6 mld-proxy
- ipv6 mld-proxy unsolicited-report-interval
- ipv6 mld-proxy version

Establishing the MLD Proxy Baseline

You can set the counters for the numbers of queries received and reports sent on the upstream interface to zero. This feature allows you to establish a reference point for MLD proxy statistics.

To set the counters:

- Issue the **baseline ipv6 mld-proxy interface** command in Global Configuration mode.

  (host1)#**baseline ipv6 mld-proxy interface**

  NOTE: Issue this command only on the upstream interface. Otherwise, this command will have no effect.

Related Documentation

- MLD Proxy Overview on page 250
- Monitoring MLD Proxy Interfaces on page 270
- baseline ipv6 mld-proxy interface

CHAPTER 12

Monitoring MLD and MLD Proxy

To display Multicast Listener Discovery (MLD) and MLD proxy parameters, use the **show** commands described in this section.

NOTE: The E120 and E320 routers output for monitor and **show** commands is identical to output from other E Series routers, except that the E120 and E320 routers output also includes information about the adapter identifier in the interface specifier (slot/adapter/port).

Monitoring MLD Information on a Virtual Router

Purpose Display MLD information for a virtual router.

Action To display MLD information for a virtual router:

```
host1:boston#show ipv6 mld
Routing Process MLD, Administrative state enabled, Operational state enabled
    2 total interfaces, 2 enabled, 0 disabled
  2 enabled interfaces performing graceful restart
    2 learnt groups
MLD Statistics:
```

```
Rcvd: 1 total, 0 checksum errors, 0 unknown types, 0 discards
      0 queries, 1 reports, 0 leaves
Sent: 2 total
```

Meaning Table 66 on page 256 lists the **show ipv6 mld** command output fields.

Table 66: show ipv6 mld Output Fields

Field Name	Field Description
Administrative state	Status of MLD in the software: enabled or disabled
Operational state	Status of MLD on the virtual router: enabled or disabled
total interfaces	Number of interfaces on which you started MLD
enabled	Number of interfaces on which MLD is enabled
disabled	Number of interfaces on which MLD is disabled
learnt groups	Number of multicast groups that the virtual router has discovered
MLD Statistics Rcvd	Statistics for MLD messages received: • total—Number of MLD messages received • checksum errors—Number of MLD messages received with checksum errors • unknown types—Number of messages received that are not multicast listener queries, multicast listener reports, or multicast listener done messages • discards—Number of multicast listener discards • queries—Number of multicast listener queries • reports—Number of multicast listener reports • leaves—Number of done messages
MLD Statistics Sent	Number of multicast listener queries sent

Related Documentation

- Enabling MLD on an Interface on page 237
- Configuring MLD Attributes on page 246
- Disabling and Removing MLD on page 249

- show ipv6 mld

Monitoring MLD Groups

Purpose Display information about statically joined and directly connected groups learned through MLD.

Action To display statically joined and directly connected groups learned through MLD without OIF mapping:

```
host1:boston# show ipv6 mld groups
    Grp Address            Interface          State      Reporter     ExpTim oldHTo
------------------- ---------------------- -------- --------------- ------ ------
ff0e::1             ATM2/0.15              Version2 fe80::90:1a02:1 54     0
                                                    640:91d
ff0e::4:1           ATM2/0.15              Version2 fe80::90:1a02:1 54     0
                                                    640:91d
  Included Sources:
    51::1                                                           54
    51::2                                                           54
Counts: 2 version-2, 0 version-1, 0 check state, 0 disabled
       (2 total)
        0 excluded
        Source-groups: 2 included, 0 excluded
```

To display statically joined and directly connected groups learned through MLD with OIF mapping:

```
host1:boston# show ipv6 mld groups
    Grp Address            Interface          State      Reporter     ExpTim oldHTo
------------------- ---------------------- -------- --------------- ------ ------
ff3e::1             ATM5/0.12              Version2 fe80::f7:0:91a: 377    0
                                                    0
                    oif-map OIFMAP ATM5/0.
                    121
ff3e::1             ATM5/0.13              Version2 fe80::f7:0:a1a: 369    0
                                                    0
                    oif-map OIFMAP ATM5/0.
                    121
ff3e::2             ATM5/0.12              Version2 fe80::f7:0:91a: 370    0
                                                    0
  Included Sources:
    10::2           oif-map OIFMAP self                             370
    10::10          oif-map OIFMAP ATM5/0.                          370
                    120
    10::11          oif-map OIFMAP ATM5/0.                          370
                    121
ff3e::2             ATM5/0.13              Version2 fe80::f7:0:a1a: 373    0
                                                    0
  Included Sources:
    10::2           oif-map OIFMAP self                             373
    10::10          oif-map OIFMAP ATM5/0.                          373
                    120
    10::11          oif-map OIFMAP ATM5/0.                          373
                    121

Counts: 4 version-2, 0 version-1, 0 check state, 0 disabled
       (4 total)
        0 excluded
        Source-groups: 6 included, 0 excluded
```

Meaning Table 67 on page 258 lists the **show ipv6 mld groups** command output fields.

Table 67: show ipv6 mld groups Output Fields

Field Name	Field Description
Grp Address	Address of the multicast group
Interface	Interface that discovered the multicast group
oif-map	Name of the OIF map and the mapped OIF interface, if a group or source has been mapped to an OIF
State	MLD version of the group
Reporter	Link-local address of the host reporting the multicast group
ExpTim	Remaining time, in seconds, at which the router stops polling for more members of this group
oldHTo	Remaining time at which the router stops polling for more MLDv1 members of a group. If this value is 0, the interface has received no MLDv1 reports for the group.
Included Sources	Sources included in the multicast group
Excluded Sources	Sources excluded from the multicast group
Counts	Number of source-group mappings by version and state

Related Documentation

- Enabling MLD on an Interface on page 237
- Configuring Multicast Groups for MLD on page 240
- Specifying MLD Multicast Groups on page 241
- Assigning an MLD Multicast Group to an Interface on page 242
- Configuring MLD Group Outgoing Interface Mapping on page 242
- show ipv6 mld groups

Monitoring MLD Interfaces

Purpose Display MLD information for interfaces on which you enabled MLD. You can use the **brief** keyword to see a summary of the information. You can use the **count** keyword to see the number of MLD interfaces.

Action To display MLD information for all interfaces on which you enabled MLD:

```
host1:boston#show ipv6 mld interface
```

```
Interface ATM2/0.103 address 13.0.0.1/255.255.255.0
Administrative state enabled, Operational state enabled
```



```
Interface parameters:
    Version 2
    State Querier
    Query Interval 125 secs, 125 secs before the next query
    Other querier present interval 250 secs
    Maximum response time 255 (in 10ths of a second)
    Last member query interval 10 (in 10ths of a second)
    Robustness 3
    No inbound access group
    No inbound access source-group
    No inbound apply-oif-map
    Immediate Leave: disabled
Explicit Host Tracking: enabled
Max-Group limit: No Limit
Admission-Bandwidth limit: No limit
IOA Packet Replication: None
    Group Count: 1
Interface statistics:
    Rcvd: 2 reports, 0 leaves, 0 wrong version queries
    Sent: 2 queries
    Groups learnt: 1

Counts: 0 down, 0 init state, 1 querier, 0 non-querier, 1 Total
```

To display MLD information for the interface on which you enabled MLD:

```
host1#show ipv6 mld interface gigabitEthernet 3/0.0
Interface GigabitEthernet3/0.0 address 10.1.1.1/255.255.255.0
Administrative state enabled, Operational state enabled
Interface parameters:
    Version 1
    State Querier
    Query Interval 125 secs, 123 secs before the next query
    Other querier present interval 250 secs
    Maximum response time 100 (in 10ths of a second)
    Last member query interval 10 (in 10ths of a second)
    Robustness 3
    No inbound access group
    No inbound access source-group
    No inbound apply-oif-map
    Immediate Leave: disabled
Explicit Host Tracking: enabled
Max-Group limit: No Limit
    Group Count: 0
    IOA packet replication gigabitEthernet 3/8.1
Interface statistics:
    Rcvd: 0 reports, 0 leaves, 0 wrong version queries
    Sent: 14 queries
    Groups learnt: 0

Counts: 0 down, 0 init state, 1 querier, 0 non-querier, 1 Total
```

Meaning Table 68 on page 260 lists the **show ipv6 mld interface** command output fields.

Table 68: show ipv6 mld interface Output Fields

Field Name	Field Description
Interface	Type of interface and interface specifier. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.
address	IPv6 link-local address of the interface
Administrative state	Status of the interface in the software: enabled or disabled
Operational state	Physical status of the interface: enabled or disabled
Version	MLD version
State	Function of the interface: querier or nonquerier
Query Interval	Time interval at which this interface sends query messages
Other querier present interval	Time that the interface waits before declaring itself as the querier
Maximum response time	Time interval during which this interface expects a host to respond
Graceful restart	Status of graceful restart: active or complete
Last member query interval	Time that this interface waits before sending a new query to a host that sends a group leave message
Robustness	Number of times this interface sends MLD messages
Inbound access group	Information about IPv6 access lists configured with the **ipv6 mld access-group** command
No inbound access group	No IPv6 access list configured with the **ipv6 mld access-group** command
Inbound access source-group	Information about IPv6 access lists configured with the **ipv6 mld access-source-group** command
No inbound access source-group	No IPv6 access list configured with the **ipv6 mld access-source-group** command
Inbound apply-oif-map	Information about OIF map names configured with the **ipv6 mld apply-oif-map** command
No inbound apply-oif-map	No OIF map name configured with the **ipv6 mld apply-oif-map** command

Table 68: show ipv6 mld interface Output Fields *(continued)*

Field Name	Field Description
Immediate Leave	Setting of the **ipv6 mld immediate-leave** command: enabled or disabled
Explicit Host Tracking	Setting of the **ipv6 mld explicit-tracking** command: enabled or disabled
Max-Group limit	Number of MLD groups that the interface can accept, as configured with the **ipv6 mld group limit** command
Group Count	Number of MLD groups that the interface has accepted
IOA packet replication	Hardware multicast packet replication interface to which egress multicast packets on this interface are redirected
Interface statistics Rcvd	Information about MLD messages received on this interface: • reports—Number of group multicast listener reports received • leaves—Number of group multicast listener done messages received • wrong version queries—Number of multicast listener queries received from devices running a different version of MLD
Interface statistics Sent	Number of MLD messages this interface has sent
Interface statistics Groups learnt	Number of groups this interface has discovered
Counts	Total number of MLD interfaces

Related Documentation

- Enabling MLD on an Interface on page 237
- Configuring MLD Settings on an Interface on page 238
- Configuring Multicast Groups for MLD on page 240
- Specifying MLD Multicast Groups on page 241
- Configuring MLD Group Outgoing Interface Mapping on page 242
- Configuring MLD Attributes on page 246
- Limiting the Number of Accepted MLD Groups on page 247
- Configuring MLD Explicit Host Tracking on page 248

- ipv6 mld access-group
- ipv6 mld access-source-group

- ipv6 mld apply-oif-map
- ipv6 mld explicit-tracking
- ipv6 mld group limit
- ipv6 mld immediate-leave
- show ipv6 mld interface

Monitoring Summary Information for MLD Interfaces

Purpose Display a summary of MLD information for interfaces on which you enabled MLD.

Action To display a summary of MLD information for interfaces on which you enabled MLD:

```
host1:boston# show ipv6 mld interface brief
   Interface      Intf Address    Ver    State        Querier    QTime  QPTime
--------------- --------------- --- ---------- ------------- ------ ------
ATM2/0.103      130.0.0.1/24    2   Querier     130.0.0.1       3(7)    0

ATM5/0.102      131.0.0.1/24    2   Querier     130.0.0.1     101(7)    0

Counts: 0 down, 0 init state, 2 querier, 0 non-querier, 2 Total
```

Meaning Table 69 on page 262 lists the **show ipv6 mld interface brief** command output fields.

Table 69: show ipv6 mld interface brief Output Fields

Field Name	Field Description
Interface	Type of interface and interface specifier. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.
Intf Address	IPv6 link-local address of the interface
Ver	MLD version
State	Function of the interface: querier or nonquerier
Querier	IPv6 address of the querier on the network to which this interface connects
QTime	Remaining time interval at which this interface sends query messages
QPTime	Remaining time that the interface waits before declaring itself as the querier

Related Documentation

- Configuring MLD Settings on an Interface on page 238
- show ipv6 mld interface

Monitoring MLD on Mapped Outgoing Interfaces

Purpose Display the current mappings to all mapped outgoing interfaces or to the specified mapped outgoing interface.

Action To display the current mappings to all mapped outgoing interfaces:

host1# **show ipv6 mld mapped-oif**

```
      OIF       Oper  Group Address  Source Address    Join I/F      Map Name
--------------- ---- --------------- --------------- --------------- ----------
ATM5/0.120      Up   ff3e::2         10::10          ATM5/0.12       OIFMAP
                                                     ATM5/0.13       OIFMAP
ATM5/0.121      Up   ff3e::1         *               ATM5/0.12       OIFMAP
                                                     ATM5/0.13       OIFMAP
                     ff3e::2         10::11          ATM5/0.12       OIFMAP
                                                     ATM5/0.13       OIFMAP

Counts: 3 source-group mappings
```

Meaning Table 70 on page 263 lists the **show ipv6 mld mapped-oif** command output fields.

Table 70: show ipv6 mld mapped-oif Output Fields

Field Name	Field Description
OIF	OIF used in an OIF map
Oper	Operation status of the outgoing interface
Group Address	Multicast group IP address associated with the OIF
Source Address	Source IP address associated with the OIF
Join I/F	MLD protocol interface associated with the OIF
Map Name	Name of the map associated to the OIF
Counts	Number of source-group mappings to OIFs

Related Documentation

- Configuring Multicast Groups for MLD on page 240
- Configuring MLD Group Outgoing Interface Mapping on page 242
- show ipv6 mld mapped-oif

Monitoring MLD on Outgoing Interfaces

Purpose Display all OIF maps or the OIF map for the specified map name.

Action To display all OIF maps:


```
host1#show ipv6 mld oif-map
     Map Name            Group Prefix       Source Prefix            OIF
------------------ ------------------ ------------------ ------------------
OIFMAP             ff3e::/112         ::/0               ATM5/0.121
                   ff3e::/112         10::2/128          self
                   ff3e::/112         10::10/128         ATM5/0.120
                   ff3e::/128         ::/0               ATM5/0.130
                   ff3e::4/128        ::/0               ATM5/0.130
```

Meaning Table 71 on page 264 lists the **show ipv6 mld oif-map** command output fields.

Table 71: show ipv6 mld oif-map Output Fields

Field Name	Field Description
Map Name	Name of the map associated to the show output
Group Prefix	Multicast group IPv6 prefix
Source Prefix	Source IPv6 prefix
OIF	Outgoing interface associated with the group and source prefix

Related Documentation

- Configuring Multicast Groups for MLD on page 240
- Configuring MLD Group Outgoing Interface Mapping on page 242
- show ipv6 mld oif-map

Monitoring MLD Membership for Multicast Groups

Purpose Display MLD membership information for multicast groups and (S, G) channels. You can use the **brief** keyword to see a summary of the information. You can use the **tracked** keyword to see interface information only for interfaces where explicit host tracking is enabled.

Action To display MLD membership information for multicast groups and (S, G) channels:

```
host1:boston# show ipv6 mld membership
Flags: M - Uses Oifmap  S- SSM mapped  T - tracked
       1,2 - The version of MLD the group is in
Reporter:
<ip-address> - last reporter if the group is not explicitly tracked
<n>/<m> - <n> reporters include mode, <m> reporters in exclude
Group      Source      Reporter                 ExpTim    Flags     Interface
---------  ---------   ----------------------   ------    ------    ---------------
ff0e::40     *         fe80::90:1a02:1640:91d   02:41     2S        FastEthernet2/1

ff0e::50               1/2                      02:56    3MT       FastEthernet2/2
                       fe80::90:1a02:1640:911   02:30
                       fe80::90:1a02:1640:912   02:48
           20::11
                       fe80::90:1a02:1640:913   02:56
           20::12
```



```
                        fe80::90:1a02:1640:911    02:30
           20::13
                        fe80::90:1a02:1640:911    02:30
                        fe80::90:1a02:1640:912    02:48
                        fe80::90:1a02:1640:913    02:56
ff0e::60                fe80::90:1a02:1640:901    01:56     3        FastEthernet2/3
           10::10                                 02:45
           10::11                                 02:35
           10::12                                 02:15
           10::14                                 stop
ff0e::70                fe80::90:1a02:1640:91     stop      3        FastEthernet2/4
           40::10                                 01:10
           40::11                                 01.24
ff0e::80                2/0                       stop      3T       FastEthernet2/5
           50::10
                        fe80::90:1a02:1650:910    02:48
           50::11
                        fe80::90:1a02:1650:920    02:56
                        fe80::90:1a02:1650:910    02:48
           50::12       fe80::90:1a02:1650:920    02:56
ff0e::90                0/3                       02:56     2T       FastEthernet2/6
              *
                        fe80::90:1a02:1660:910    02:48
                        fe80::90:1a02:1660:920    02:56
                        fe80::90:1a02:1660:930    02:48
```

Meaning Table 72 on page 265 lists the **show ipv6 mld membership** command output fields.

Table 72: show ipv6 mld membership Output Fields

Field Name	Field Description
Group	Multicast group or (S, G) channel
Source	(S, G) entries that are forwarding traffic
Reporter	Hosts that requested including sources or that have not requested excluding sources. If listed under a group, host that sent exclude reports for the group. If listed under a source, host that requested traffic from this source for the group. For any (S, G), if listed under a source, indicates hosts interested in the traffic for this (S, G).
ExpTim	Expiration time
Flags	• M—Uses Oifmap • S—SSM mapped • T—Tracked • 1, 2—MLD version that the group is in
Interface	Type of interface and interface specifier. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.

Related Documentation

- Configuring Multicast Groups for MLD on page 240
- Assigning an MLD Multicast Group to an Interface on page 242
- Configuring MLD Attributes on page 246
- Including and Excluding MLD Traffic on page 248
- Configuring MLD Explicit Host Tracking on page 248
- show ipv6 mld membership

Monitoring MLD Information for Mapped Outgoing Interfaces

Purpose Display the OIF to be assigned to a given map-name, group address, and source address.

Action To display the mapped OIF to be assigned to a given map-name, group address, and source address:

```
host1#show ipv6 mld oif-mapping OIFMAP ff3e::1 10::10
    OIF Mapping
    OIF-MAP Name   : OIFMAP
    Group Address  : ff3e::1
    Source Address : 10::10
    Mapped OIF     : ATM5/0.120
```

Meaning Table 73 on page 266 lists the **show ipv6 mld oif-mapping** command output fields.

Table 73: show ipv6 mld oif-mapping Output Fields

Field Name	Field Description
OIF-MAP Name	Name of the map requested
Group Address	Multicast group IP address requested
Source Address	Source IP address requested
Mapped OIF	Join interface associated with the OIF map

Related Documentation

- Configuring Multicast Groups for MLD on page 240
- Configuring MLD Group Outgoing Interface Mapping on page 242
- show ipv6 mld oif-mapping

Monitoring MLD SSM Mapping

Purpose Display the SSM mapping state and the source list mapping associated with a multicast group address, based on the static SSM mapping configuration.

Action To display the SSM mapping state and the source list mapping associated with a multicast group address:


```
host1:boston#show ipv6 mld ssm-mapping ff3e::1
    SSM Mapping   : Enabled
    Group Address : ff3e::1
    Source List   : 2001::1
                  : 2001::2
```

Meaning Table 74 on page 267 lists the **show ipv6 mld ssm-mapping** command output fields.

Table 74: show ipv6 mld ssm-mapping Output Fields

Field Name	Field Description
SSM Mapping	Status of SSM mapping on the interface (enabled or disabled)
Group Address	Multicast group address requested
Source List	List of sources mapped to the multicast group address

Related Documentation

- Configuring MLD Attributes on page 246
- Configuring MLD SSM Mapping on page 247
- show ipv6 mld ssm-mapping

Monitoring the Number of MLD Groups on a Port

Purpose Display the number of MLD groups that ports have accepted and, if configured, the maximum number of groups that ports can accept. A value of –1 indicates that no port group limit is configured. Only ports that have accepted MLD groups and ports for which you have configured a limit for the number of MLD groups appear in this display.

Action To display the number of MLD groups that ports have accepted and the maximum number of groups that ports can accept:

```
host1:boston#show multicast group limit
Port      limit count
--------- ----- -----
2/0           5     0
2/1          -1     1
```

Meaning Table 75 on page 267 lists the **show multicast group limit** command output fields.

Table 75: show multicast group limit Output Fields

Field Name	Field Description
Port	Identifier of the port in *slot/port* format: • slot—Number of the chassis slot in the range 0–6 (ERX7xx models) or 0–13 (ERX14xx models) • port—Port number on the I/O module

Table 75: show multicast group limit Output Fields *(continued)*

Field Name	Field Description
limit	Maximum number of MLD groups that the port can accept. A value of –1 indicates that no limit has been specified.
count	Actual number of MLD groups the port has accepted

Related Documentation

- Configuring MLD Attributes on page 246
- Limiting the Number of Accepted MLD Groups on page 247
- show multicast group limit

Monitoring MLD Proxy Parameters

Purpose Display MLD proxy parameters for a virtual router.

Action To display MLD proxy parameters for a virtual router:

```
host1#sshow ipv6 mld-proxy
Routing Process MLD Proxy, Administrative state enabled, Operational state
enabled
    total 1 upstream interface, state enabled
    1 multicast group
```

Meaning Table 76 on page 268 lists the **show ipv6 mld-proxy** command output fields.

Table 76: show ipv6 mld-proxy Output Fields

Field Name	Field Description
Routing Process	MLD proxy protocol
Administrative state	State of MLD proxy in the software
Operational state	Operational state of MLD proxy: enabled or disabled
total interfaces	Number of MLD proxy interfaces on the virtual router; currently only one upstream interface per virtual router
state	Operational state of the MLD proxy interfaces: enabled or disabled
multicast group	Number of multicast groups associated with MLD proxy interfaces

Related Documentation

- Configuring MLD Proxy on page 251
- show ipv6 mld-proxy

Monitoring MLD Proxy Groups

Purpose Display information about all or specified multicast groups that MLD proxy reported.

Action To display information about all multicast groups that MLD proxy reported:

host1#**show ipv6 mld-proxy groups**

```
Grp Address     Interface       Grp Mode
--------------- --------------- --------------
ff0e::1         ATM5/1.200
ff0e::2         ATM5/1.200
ff0e::3         ATM5/1.200      Include(1):
  2001::1
ff0e::4         ATM5/1.200
ff0e::5         ATM5/1.200      Exclude(1):
  2001::2

Counts: 3 <*,G>, 1 Exclude (1 sources), 1 Include (1 sources)
       (5 total)
```

To display information about specified multicast groups that MLD proxy reported:

host1#**show ipv6 mld-proxy groups ff0e::1**

```
Grp Address     Interface       Grp Mode
--------------- --------------- --------------
ff0e::1         ATM5/1.200

Counts: 1 <*,G>
       (1 total)
```

To display the number of multicast groups that MLD proxy reported:

host1#**show ipv6 mld-proxy groups count**

```
Counts: 3 <*,G>, 1 Exclude (1 sources), 1 Include (1 sources)
       (5 total)
```

Meaning Table 77 on page 269 lists the **show ipv6 mld-proxy groups** command output fields.

Table 77: show ipv6 mld-proxy groups Output Fields

Field Name	Field Description
Grp Address	Address of the multicast group
Interface	Type and identifier of the upstream interface associated with the multicast group
Grp Mode	• Blank—No sources included or excluded for this group • Include—Sources included for this group • Exclude—Sources excluded for this group
Count	Total number of multicast groups associated with this interface

Related Documentation

- show ipv6 mld-proxy groups

Monitoring MLD Proxy Interfaces

Purpose

Display information about all or specified interfaces on which you configured MLD proxy. To view information about a particular interface, enter an interface type and specifier, such as atm 3/0. You can use the **brief** option to display a summary rather than a detailed description.

Action

To display information about all interfaces on which you configured MLD proxy:

`host1#`**show ipv6 mld-proxy interface**

```
Interface ATM5/1.200 address fe80::f7:0:231a:0
Administrative state enabled, Operational state enabled
Interface parameters:
    Version 2
    State No v1 Router Present
    Unsolicited report interval 100 (in 10ths of a second)
    5 multicast groups
Interface statistics:
    Rcvd: 0 v1 query, 0 v1 report, 25 v2 queries, 0 v2 report
    Sent: 0 v1 report, 0 v1 leave, 35 v2 reports
```

To display summarized information about all interfaces on which you configured MLD proxy:

`host1#`**show ipv6 mld-proxy interface brief**

```
Interface       Intf Address       Ver State                UnSlTime
--------------- ------------------ --- -------------------- --------
ATM5/1.200      fe80::f7:0:231a:0    2 No v1 Router Present      100
```

Meaning

Table 78 on page 270 lists the **show ipv6 mld-proxy interface** command output fields.

Table 78: show ipv6 mld-proxy interface Output Fields

Field Name	Field Description
Interface	Type of upstream interface. For details about interface types, see Interface Types and Specifiers in *JunosE Command Reference Guide*.
Address	Address of upstream interface
Administrative state	State of upstream interface in the software: enabled or disabled
Operational state	Physical state of upstream interface: enabled or disabled
Version	MLD version on this interface
State	Presence of MLDv1 routers on the same subnetwork as this upstream interface

Table 78: show ipv6 mld-proxy interface Output Fields *(continued)*

Field Name	Field Description
Unsolicited report interval	Time interval at which this upstream interface sends unsolicited group membership report
multicast group	Number of multicast groups associated with this upstream interface
Interface statistics Rcvd	Statistics for messages received on this interface: • v1 queries—Number of MLDv1 multicast listener queries received • v1 report—Number of MLDv1 multicast listener reports received • v2 queries—Number of MLDv2 multicast listener queries received • v2 report—Number of MLDv2 multicast listener reports received
Interface statistics Sent	Statistics for messages sent from this interface: • v1 reports—Number of MLDv1 multicast listener reports sent • v1 leaves—Number of multicast listener done messages sent • v2 reports—Number of MLDv2 multicast listener reports sent

Related Documentation

- Configuring MLD Proxy on page 251
- Establishing the MLD Proxy Baseline on page 252
- show ipv6 mld-proxy interface

CHAPTER 13

Configuring PIM for IPv6 Multicast

Protocol Independent Multicast (PIM) is a collection of multicast routing protocols that enable multicast routers to identify other multicast routers that can receive packets.

This chapter describes how to configure PIM for IPv6 multicast on the E Series router; it contains the following sections:

Understanding PIM for IPv6 Multicast

This implementation of PIM supports PIM sparse mode and PIM source-specific multicast (PIM SSM) for IPv6 multicast.

SSM is an extension to the Any Source Multicast (ASM) service model and facilitates the deployment of broadcast (one-to-many) applications, such as Internet TV and radio where large receiver audiences require traffic from a few well-known sources.

Figure 21 on page 274 represents how PIM builds an (S,G) entry in an SRT. When multiple routers are connected to a multiaccess network, one router is assigned the role of the designated router. The designated router receives data from the source on interface 1/0 and multicasts the data to its downstream neighbors on interfaces 1/1, 2/0, and 2/1. In the designated router routing table, the entry for this operation lists the source as the IP address of the source and the group as the IP address of the multicast group.

Neighbors exchange hello messages periodically to determine the designated router. The router with the highest network layer address becomes the designated router. If the designated router subsequently receives a hello message from a neighbor with a higher network layer address, that neighbor becomes the designated router.

Figure 21: Source-Rooted Tree

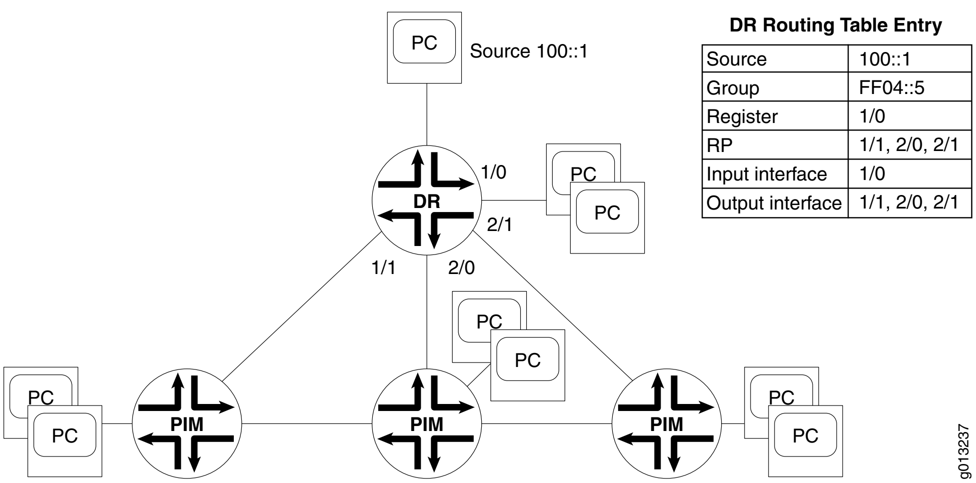

The IPv6 implementation of PIM supports the following modes:

- PIM Sparse Mode on page 274
- PIM Sparse Mode Bootstrap Router on page 275
- PIM Source-Specific Multicast on page 275

PIM Sparse Mode

In addition to the features PIM sparse mode supports for IPv4, this IPv6 implementation of PIM sparse mode also supports remote neighbors.

For a description of PIM sparse mode, see "Configuring PIM for IPv4 Multicast" on page 89.

Joining Groups

A host's designated router (DR) sends join messages to the RP when that host wants to join a group. When a host wants to leave a group, it communicates with its designated router through MLD. When the designated router no longer has any hosts that belong to a particular group, it sends a prune message to the RP.

Timers

PIM sparse mode uses timers to maintain the networking trees.

NOTE: **PIM sparse mode routers poll their neighbors and hosts for various pieces of information at set intervals.**

If a PIM sparse mode router does not receive information from a neighbor or host within a specific time, known as the *hold time*, it removes the associated information from its routing tables.

You can configure how often an interface sends hello messages (hello interval) and how often routers send RP announce messages (RP announce interval). The hold-time associated with hello messages is 3.5 times the hello interval, and the holdtime associated with RP announce messages is 2.5 times the RP announce interval.

All other timers are fixed and take the default values recommended in:

RFC 2934—Protocol Independent Multicast MIB for IPv4 (October 2000)

PIM Sparse Mode Bootstrap Router

PIM sparse mode routers need the address of the rendezvous point (RP) for each group for which they have (*,G) state. They obtain this address either through a bootstrap mechanism or through static configuration. Two bootstrap mechanisms exist—bootstrap router (BSR) or auto-RP. Auto-RP is not used in IPv6 implementations.

When implemented, BSR operates as follows:

1. One router in each PIM domain is elected the BSR.
2. All the routers in the domain that are configured to be RP candidates periodically unicast their candidacy to the BSR.
3. The BSR picks an RP set from the available candidates and periodically announces this set in a bootstrap message.
4. Bootstrap messages are flooded hop by hop throughout the domain until all routers in the domain learn the RP Set.

PIM Source-Specific Multicast

PIM source-specific multicast (SSM) is an extension of the PIM protocol. Using SSM, a client can receive multicast traffic directly from the source. PIM SSM uses PIM sparse mode functionality to create an SPT between the client and the source, but builds the SPT without using an RP.

By default, the SSM group multicast address is limited to the IPv6 address range FF3x::/96 where x represents any valid scope. You can use the **ipv6 pim ssm range** command to change the SSM group address range.

Advantages that an SSM-configured network has over a traditionally configured PIM sparse mode network include the following:

- No need for shared trees or RP mapping (no RP is required).
- No need for RP-to-RP source discovery through Multicast Source Discovery Protocol (MSDP).
- Simplified administrative deployment; you need only configure PIM sparse mode on all router interfaces and issue the necessary SSM commands (including specifying MLDv2 on the receiver local area network).
- Support for source lists; you can use source lists, supported in MLDv2, where only specified sources send traffic to the SSM group.

In a PIM SSM–configured network, the E Series router subscribes to an SSM channel (by means of MLDv2), announcing a desire to join group G and source S. The directly connected PIM sparse mode router, the designated router of the receiver, sends an (S,G) join message to its RPF neighbor for the source. For PIM SSM, the RP is not contacted in this process by the receiver (as happens in normal PIM sparse mode operations).

NOTE: You can configure multicast on IPv4 and IPv6 interfaces. For information about configuring PIM on IPv4 interfaces, see “Configuring PIM for IPv4 Multicast” on page 89.
For information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- PIM for IPv6 Multicast Platform Considerations on page 276
- PIM for IPv6 Multicast References on page 277

PIM for IPv6 Multicast Platform Considerations

For information about modules that support PIM for IPv6 multicasting on the ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support PIM for IPv6 multicasting.

For information about modules that support PIM for IPv6 multicasting on the E120 and E320 Broadband Services Routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support PIM for IPv6 multicasting.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- PIM for IPv6 Multicast References on page 277
- Monitoring PIM for IPv6 Multicast on page 293

PIM for IPv6 Multicast References

For more information about IPv6 multicast, see the following resources:

- RFC 2362—Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification (June 1998)
- RFC 3569—An Overview of Source-Specific Multicast (SSM) (July 2003)
- Source-Specific Multicast for IP—draft-ietf-ssm-arch-06.txt (March 2005 expiration)
- Source-Specific Protocol Independent Multicast in 232/8—draft-ietf-mboned-ssm232-08.txt (September 2004 expiration)

NOTE: **IETF drafts are valid for only 6 months from the date of issuance. They must be considered as works in progress. Please refer to the IETF Web site at http://www.ietf.org for the latest drafts.**

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- PIM for IPv6 Multicast Platform Considerations on page 276
- Monitoring PIM for IPv6 Multicast on page 293

Enabling PIM for IPv6 on a Virtual Router

By default, PIM is disabled. To enable PIM on a VR:

1. Enable multicast routing.
2. Create a VR, or access the VR context.
3. Create and enable PIM processing.

```
host1(config)#virtual-router boston
host1:boston(config)#ipv6 router pim
```

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Disabling PIM for IPv6 on a Virtual Router on page 278
- Enabling PIM for IPv6 on an Interface on page 278
- Removing PIM for IPv6 on page 290
- Resetting PIM Counters and Mappings for IPv6 on page 290
- **ipv6 router pim**
- **virtual-router**

Disabling PIM for IPv6 on a Virtual Router

To disable PIM processing on a router, use the **pim disable** command. By default, PIM processing is enabled:

- Issue the **pim disable** command in Router Configuration mode.

  host1:boston(config-router)#pim disable

  Use the **no** version to reenable PIM processing.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Enabling PIM for IPv6 on a Virtual Router on page 277
- Enabling PIM for IPv6 on an Interface on page 278
- Removing PIM for IPv6 on page 290
- Resetting PIM Counters and Mappings for IPv6 on page 290
- **pim disable**

Enabling PIM for IPv6 on an Interface

You can enable PIM on an interface in one of the allowed modes and specify how often the interface sends hello messages to neighbors.

You can configure PIM and MLD on the same interface. If you configure MLD and PIM on an interface, the router considers that PIM owns the interface.

To enable PIM on an interface:

1. Enable PIM in sparse mode on an interface.

   host1(config-if)#ipv6 pim sparse-mode

   Use the **no** version to disable PIM in sparse mode on an interface.

2. Specify the interval, in seconds, at which the router sends hello messages to neighbors.

   host1(config-if)#ipv6 pim query-interval 100

   Use the **no** version to restore the default setting, 30 seconds.

3. Set the graceful restart duration for IP PIM sparse mode.

   host1(config-if)#**ip pimv6 sparse-mode graceful-restart-duration 10**

   Use the **no** version to return to the default duration of 30 seconds.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Enabling PIM for IPv6 on a Virtual Router on page 277
- Removing PIM for IPv6 on page 290
- Resetting PIM Counters and Mappings for IPv6 on page 290
- **ipv6 pim query-interval**
- **ipv6 pim sparse-mode**
- **ipv6 pim sparse-mode graceful-restart-duration**

Setting a Priority to Determine the Designated Router for IPv6

You can influence whether a particular router is selected as the designated router with the **ipv6 pim dr-priority** command. A higher priority value increases the likelihood that a router is selected as the designated router, while a lower value decreases the likelihood. The **ipv6 pim dr-priority** command in Router Configuration mode sets the designated router priority on all the PIM interfaces on the router. To override this global setting on a particular interface, use the **ipv6 pim dr-priority** command in Interface Configuration mode.

NOTE: You cannot configure the designated router priority on PIM DM interfaces.

- To set a priority value, in the range 1–254, by which a router is likely to be selected as the designated router, issue the **ip pim dr-priority** command.

  In Router Configuration mode:

  host1(config-router)#**ipv6 pim dr-priority 24**

  The **no** version restores the default value. The default value is 1.

  In Interface Configuration mode:

  host1(config-subif)#**ipv6 pim dr-priority 24**

  The **no** version restores the value that is specified in Router Configuration mode.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Enabling PIM for IPv6 on a Virtual Router on page 277
- Disabling PIM for IPv6 on a Virtual Router on page 278
- Enabling PIM for IPv6 on an Interface on page 278

- Removing PIM for IPv6 on page 290
- Monitoring PIM for IPv6 Multicast on page 293
- **ipv6 pim dr-priority**

Configuring the PIM Join/Prune Message Interval for IPv6

When you use the router for PIM, the router sends join/prune message to the upstream RPF neighbor. The default join/prune message interval is 60 seconds. You can configure the join/prune message interval using the **ipv6 pim join-prune-interval** command. The **ipv6 pim join-prune-interval** command in Router Configuration mode configures the join-prune interval on all the PIM interfaces on the router. To override this global setting on a particular interface, use the **ipv6 pim join-prune-interval** command in Interface Configuration mode.

The hold-time associated with the PIM join/prune interval messages is 3.5 times the PIM join/prune message interval.

NOTE: **You cannot configure the PIM join/prune message interval on PIM dense mode interfaces.**

- To set an interval value, in the range 10–210 seconds, at which the router sends the PIM join/prune message, issue the **ipv6 pim join-prune-interval** command.

  In Router Configuration mode:

  ```
  host1(config-router)#ipv6 pim join-prune-interval 150
  ```

  The **no** version restores the default value, 60 seconds.

  In Interface Configuration mode:

  ```
  host1(config-subif)#ipv6 pim join-prune-interval 150
  ```

  The **no** version restores the value that is specified in Router Configuration mode.

NOTE: You can configure multicast on IPv4 and IPv6 interfaces. For information about configuring PIM on IPv4 interfaces, see "Configuring PIM for IPv4 Multicast" on page 89.
For information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Enabling PIM for IPv6 on a Virtual Router on page 277
- Disabling PIM for IPv6 on a Virtual Router on page 278

- Enabling PIM for IPv6 on an Interface on page 278
- Removing PIM for IPv6 on page 290
- Monitoring PIM for IPv6 Multicast on page 293
- **ipv6 pim join-prune-interval**

Configuring an RP Router for PIM Sparse Mode for IPv6

When you use the router for PIM sparse mode, some VRs must act as RP routers. If you want to control PIM more tightly, you can configure a static RP router. To do so:

1. Configure an access list that details the multicast groups that can use the static RP router (in this case, all globally scoped multicast groups).

   host1(config)#**ipv6 access-list boston permit ff0e::/16 any**

2. Specify a static RP router.

   host1(config)#**ipv6 pim rp-address ::122:1 boston**

NOTE: You can configure multicast on IPv4 and IPv6 interfaces. For information about configuring PIM on IPv4 interfaces, see "Configuring PIM for IPv4 Multicast" on page 89.
For information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Enabling PIM for IPv6 on a Virtual Router on page 277
- Disabling PIM for IPv6 on a Virtual Router on page 278
- Enabling PIM for IPv6 on an Interface on page 278
- Removing PIM for IPv6 on page 290
- Monitoring PIM for IPv6 Multicast on page 293
- **ipv6 pim rp-address**

Configuring BSR and RP Candidates for PIM Sparse Mode for IPv6

When choosing candidate BSRs or candidate RPs, select well-connected routers in the core of the network.

Typically, candidate BSRs are a subset of the candidate RPs. A single BSR is elected for the domain the set of candidate BSRs. The elected BSR floods bootstrap messages

(BSMs) containing their group-to-RP mappings to all PIM routers. PIM routers use the group-to-RP mappings supplied by the elected (or preferred) BSR.

Candidate RPs are routers that are capable of performing as a rendezvous point router for one or more multicast groups. Candidate RPs periodically advertise the set of groups they support to BSRs. A candidate RP may support all the multicast group address range or any subset thereof. You can achieve redundancy by configuring more than one candidate RP for a group or range of groups.

- Issue the **ipv6 pim bsr-candidate** command in Global Configuration mode to define a router as a BSR candidate:

  ```
  host1(config)#ipv6 pim bsr-candidate loopback 1 30 10
  ```

  The **no** version stop the router from acting as a BSR candidate.

- Issue the **ipv6 pim rp-candidate** command in Global Configuration mode to define a router as an RP candidate:

  ```
  host1(config)#ipv6 access-list 1 permit  1001::1
  host1(config)#ipv6 access-list 1 permit  1002::1
  host1(config)#ipv6 pim rp-candidate loopback 1 group-list 1
  ```

  The **no** version stops the router from being an RP candidate.

NOTE: You can configure multicast on IPv4 and IPv6 interfaces. For information about configuring PIM on IPv4 interfaces, see "Configuring PIM for IPv4 Multicast" on page 89.
For information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Monitoring PIM for IPv6 Multicast on page 293
- **ipv6 pim bsr-candidate**
- **ipv6 pim rp-candidate**

Switching to an SPT for PIM Sparse Mode for IPv6

PIM sparse mode initiates multicast using a shared tree. You can configure PIM sparse mode to switch to an SPT when a source starts sending multicast messages, or you can prevent PIM sparse mode from switching to an SPT. Multicasting over an SPT can be more efficient than multicasting over a shared tree (see *PIM Sparse Mode* in "Understanding PIM for IPv6 Multicast" on page 274).

To specify when PIM sparse mode switches from a shared tree to an SPT:

- Issue the **ipv6 pim spt-threshold** command in Global Configuration mode.

host1(config)#**ipv6 pim spt-threshold 4**

The **no** version restores the default value, 0.

NOTE: You can configure multicast on IPv4 and IPv6 interfaces. For information about configuring PIM on IPv4 interfaces, see "Configuring PIM for IPv4 Multicast" on page 89.
For information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Monitoring PIM for IPv6 Multicast on page 293

- **ipv6 pim spt-threshold**

Configuring PIM Sparse Mode Remote Neighbors for IPv6

You must use PIM sparse mode remote neighbors to run multicast services over BGP/MPLS VPNs.

NOTE: Although you can configure PIM sparse mode remote neighbors, you cannot use these remote neighbors for BGP/MPLS VPNs.

To configure a pair of E Series routers to act as PIM remote neighbors:

1. On one router, specify the other router to be a remote neighbor, and identify the IP address of the interface on the other router that is used for the connection to this router.

   host1(config-router):boston#**remote-neighbor 1001::1 sparse-mode**

2. Specify the location of the local interface whose address is used as the source address for the PIM connection to a remote neighbor.

   host1(config-router-rn):boston#**update-source atm 2/1.108**

3. (Optional) Specify how often the router sends hello messages to the remote neighbor.

   host1(config-router-rn):boston#**query-interval 40**

4. Repeat Steps 2 to 3 for the other router.

NOTE: You can configure multicast on IPv4 and IPv6 interfaces. For information about configuring PIM on IPv4 interfaces, see "Configuring PIM for IPv4 Multicast" on page 89.
For information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Example: Configuring PIM Sparse Mode Remote Neighbors for IPv6 on page 284
- Monitoring PIM for IPv6 Multicast on page 293

- **query-interval**
- **remote-neighbor**
- **update-source**

Example: Configuring PIM Sparse Mode Remote Neighbors for IPv6

The following example illustrates how to configure PIM sparse mode remote neighbors for IPv6.

- Requirements on page 284
- Overview on page 284
- Configuring PIM Sparse Mode Remote Neighbors for IPv6 on page 285

Requirements

This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

This example uses the configuration shown in Figure 21 on page 274. Two E Series routers called router Boston and router Chicago are running PIM and are connected by MPLS tunnels.

Figure 22: Source-Rooted Tree

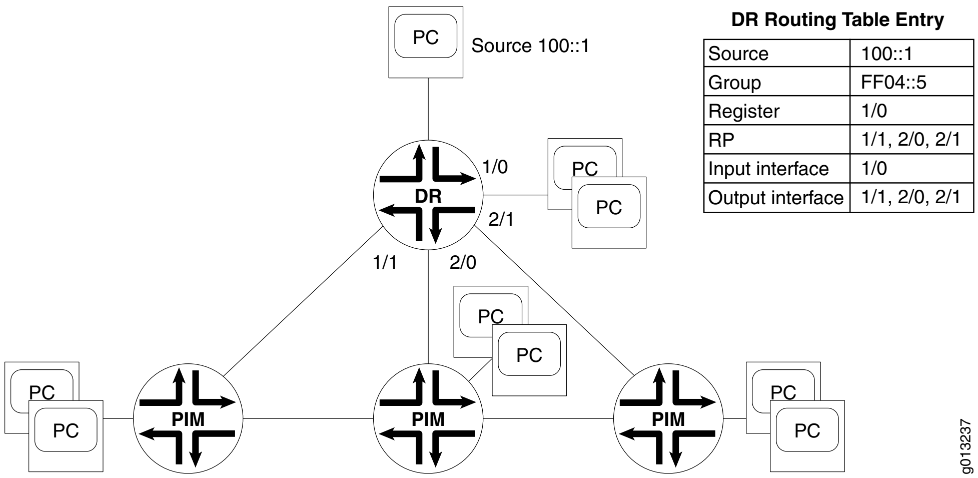

Configuring PIM Sparse Mode Remote Neighbors for IPv6

Step-by-Step Procedure

To configure the routers as PIM remote neighbors:

1. Specify that router Chicago will be a remote neighbor of router Boston, and identify the IP address on router Chicago that will transmit datagrams to router Boston.

   boston(config-router)#remote-neighbor 1001::1 sparse-mode

2. Specify the location of the interface that will transmit datagrams from router Boston to router Chicago.

   boston(config-router-rn)#update-source atm 2/1.108

3. Specify that router Boston will send hello messages to router Chicago every 40 seconds.

   boston(config-if)#ipv6 pim query-interval 40

4. Specify that router Boston will be a remote neighbor of router Chicago, and identify the IP address on router Boston that will transmit datagrams to router Chicago.

   chicago(config-router)#remote-neighbor 2001::1 sparse-mode

5. Specify the location of the interface that will transmit datagrams from router Chicago to router Boston.

   chicago(config-router-rn)#update-source atm 2/1.95

6. Specify that router Chicago will send hello messages to router Boston every 40 seconds.

   chicago(config-if)#ipv6 pim query-interval 40

7. See update-source.

NOTE: You can configure multicast on IPv4 and IPv6 interfaces. For information about configuring PIM on IPv4 interfaces, see "Configuring PIM for IPv4 Multicast" on page 89.

For information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Configuring PIM Sparse Mode Remote Neighbors for IPv6 on page 283
- Monitoring PIM for IPv6 Multicast on page 293

- **ipv6 pim query-interval**
- **remote-neighbor**
- **update-source**

Configuring PIM Sparse Mode Join Filters for IPv6

You can use PIM sparse mode join filters to prevent multicast state from being created in the PIM sparse mode router. The filters are applied to join entries in PIM join/prune messages that are received from PIM sparse mode neighbors.

By denying joins at the edge of a network, you can limit the multicast state and traffic in the network. By accepting only certain joins, you can control which multicast services an end user can receive. PIM join filters also reduce the potential for denial of service (DOS) attacks where large numbers of joins forwarded to each router on the RPT can result in a PIM state explosion and very high memory consumption.

For information about how to create access lists, see *Configuring Routing Policy* in the *JunosE IP Services Configuration Guide*.

To specify an extended access list that you want this PIM interface to use as a join filter:

- Issue the **ipv6 pim join-filter** command

  In Global Configuration mode:

  ```
  host1(config)#ipv6 pim join-filter gold
  ```

  In Interface Configuration mode:

  ```
  host1(config-interface)#ipv6 pim join-filter gold
  ```

  The **no** version removes the filter association. You can apply the join filter at the global level or at the interface level. If an interface-level filter exists, it takes precedence over the global-level filter.

NOTE: You can configure multicast on IPv4 and IPv6 interfaces. For information about configuring PIM on IPv4 interfaces, see "Configuring PIM for IPv4 Multicast" on page 89.
For information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Configuring the PIM Join/Prune Message Interval for IPv6 on page 280
- Configuring an RP Router for PIM Sparse Mode for IPv6 on page 281
- Configuring PIM Sparse Mode Remote Neighbors for IPv6 on page 283
- Monitoring PIM for IPv6 Multicast on page 293

- **ipv6 pim join-filter**

Configuring PIM for IPv6 SSM

Source Specific Multicast (SSM) is a datagram delivery model that best supports one-to-many applications, also known as broadcast applications. SSM is networking technology that targets audio and video broadcast application environments.

To use PIM SSM, MLDv2 must be configured on customer premises equipment (CPE)–facing interfaces to receivers, and PIM sparse mode must be configured on CPE-facing interfaces to sources and on core-facing interfaces. After configuring SSM, you can use the **show ipv6 pim sparse-mode sg-state** command to display SSM group membership information.

To configure PIM SSM, you enable PIM SSM on the router and define the SSM range of IP multicast addresses.

Figure 23 on page 287 shows how PIM SSM is configured between a receiver and a source in the network. Interface 1 has MLDv2 enabled and all other interfaces towards the core or source have PIM SSM enabled.

Figure 23: Network on Which to Configure PIM SSM

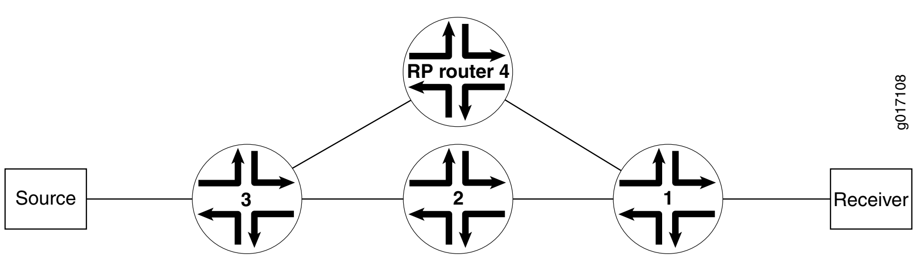

To configure PIM SSM:

1. Enable PIM SSM on the E Series router. The IANA SSM range is configured by default. You can modify the SSM address range by using the access list.

   host1(config)#**access-list 15 permit ip any host 239.0.0.2**
   host1(config)#**access-list 15 permit ip any 232.0.0.0 0.225.225.225**
   host1(config)#**ipv6 pim ssm range 15**

2. Enable PIM sparse mode on the CPE-facing interface towards the source or core.

3. Enable MLDv2 on the CPE-facing interface towards the receiver.

PIM SSM also works with MLDv1 if you configure the ssm-map in MLD as in the following example:

host1(config)#**ipv6 pim ssm**
host1(config)#**ipv6 access-list ssm_map1 permit any host ff3e::1**
host1(config)#**ipv6 mld ssm-map enable**
host1(config)#**ipv6 mld ssm-map static ssm_map1 51::1**

The **no** version disables ssm-map:

host1(config)#**no ipv6 mld ssm-map static ssm_map1 51::1**

You can enable PIM SSM and define the SSM range of IPv6 multicast addresses as in the following examples:

Example 1—Enables SSM with addresses in the IANA range. The SSM address range is set as the default, which is limited to the IPv6 address range, and where *x* represents any valid scope.

host1(config)#**ipv6 pim ssm FF3 x ::/96**

Example 2—Configures Class D addresses outside of the default SSM range.

host1(config)#**ipv6 access-list alist permit any ff3e::1:0:0/96**
host1(config)#**ipv6 pim ssm range alist**

Example 3—Resets the SSM address range to the default.

host1(config)#**ipv6 pim ssm default**

NOTE: You can configure multicast on IPv4 and IPv6 interfaces. For information about configuring PIM on IPv4 interfaces, see "Configuring PIM for IPv4 Multicast" on page 89.
For information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Monitoring PIM for IPv6 Multicast on page 293

- **access-list**
- **ipv6 pim ssm**

BFD Protocol for PIM for IPv6 Overview

The **ipv6 pim bfd-liveness-detection** command configures the Bidirectional Forwarding Detection (BFD) protocol for PIM. The BFD protocol uses control packets and shorter detection time limits to more rapidly detect failures in a network. Also, because they are adjustable, you can modify the BFD timers for more or less aggressive failure detection.

PIM routers send periodic hello messages from each PIM-enabled interface. You can configure this interval using the **ipv6 pim query-interval** command. By default, the PIM router sends a hello message every 30 seconds (with an interval range of 0–210 seconds). If it receives no response from a neighbor within 3.5 times the interval value (a minimum of 3.5 seconds), the PIM router drops the neighbor.

In contrast, when a BFD session exists between neighbors, a PIM neighbor that goes down is detected quickly (in milliseconds rather than in seconds).

When you issue the **ipv6 pim bfd-liveness-detection** command on a PIM router, the router establishes BFD liveness detection with all BFD-enabled PIM neighbors. When the local router receives an update from a remote PIM neighbor—if BFD is enabled and if the session is not already present—the local router attempts to create a BFD session to the remote neighbor.

Each adjacent pair of neighbors negotiates an acceptable transmit interval for BFD packets. The negotiated value can be different on each neighbor. Each neighbor then calculates a BFD liveness detection interval. When a neighbor does not receive a BFD packet within the detection interval, it declares the BFD session to be down.

NOTE: Before the router can use the **ipv6 pim bfd-liveness-detection** command, you must specify a BFD license key. To view an already configured license, use the **show license bfd** command.

For general information about configuring and monitoring the BFD protocol, see *Configuring BFD* in the *JunosE IP Services Configuration Guide*.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Configuring BFD Protocol for PIM for IPv6 on page 289
- Monitoring PIM for IPv6 Multicast on page 293
- **ipv6 pim bfd-liveness-detection**

Configuring BFD Protocol for PIM for IPv6

To enable BFD (bidirectional forwarding detection) and define BFD values to more quickly detect PIM data path failures:

- Issue the **ipv6 pim bfd-liveness-detection** command in Interface Configuration mode.

```
host1(config-if)#ipv6 pim bfd-liveness-detection minimum-interval 800
```

The **no** version disables BFD on the PIM interface.

NOTE: Before the router can use the **ipv6 pim bfd-liveness-detection** command, you must specify a BFD license key. To view an already configured license, use the **show license bfd** command.

For details on liveness detection negotiation, see *Negotiation of the BFD Liveness Detection Interval* in the *JunosE IP Services Configuration Guide*.

NOTE: You can configure multicast on IPv4 and IPv6 interfaces. For information about configuring PIM on IPv4 interfaces, see "Configuring PIM for IPv4 Multicast" on page 89.
For information about configuring IP interfaces, see *Configuring IP* in the *JunosE IP, IPv6, and IGP Configuration Guide*. For information about configuring IPv6 interfaces, see *Configuring IPv6* in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- BFD Protocol for PIM for IPv6 Overview on page 289
- Monitoring PIM for IPv6 Multicast on page 293

- ipv6 pim bfd-liveness-detection

Removing PIM for IPv6

To remove PIM from the VR:

- Issue the **no ipv6 router pim** command in Global Configuration mode.

```
host1:boston(config)#no ipv6 router pim
```

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Enabling PIM for IPv6 on a Virtual Router on page 277
- Disabling PIM for IPv6 on a Virtual Router on page 278
- Enabling PIM for IPv6 on an Interface on page 278
- Resetting PIM Counters and Mappings for IPv6 on page 290

- **ipv6 router pim**

Resetting PIM Counters and Mappings for IPv6

You can use the **clear ipv6 pim** commands to reset PIM counters and mappings.

- Issue the **clear ipv6 pim interface** command in Privileged Exec mode to clear the counters for multicast packet statistics on all interfaces or a specified interface:

  host1#**clear ipv6 pim interface atm 3/0.5 count**

  There is no **no** version.

  NOTE: Specify an interface type and identifier, such as atm 3/0 to clear the counters on that interface. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*. If you do not specify an interface, the router clears the counters on all interfaces.

- Issue the **clear ipv6 pim remote-neighbor** command in Privileged Exec mode to clear the counters for remote neighbor statistics on all interfaces or the specified interface:

  host1#**clear ipv6 pim remote-neighbor 1001::1 count**

  There is no **no** version.

  NOTE: Specify the IP address of an interface to clear the counters for that interface. If you do not specify an interface, the router clears the counters on all interfaces.

Related Documentation

- Understanding PIM for IPv6 Multicast on page 274
- Enabling PIM for IPv6 on a Virtual Router on page 277
- Disabling PIM for IPv6 on a Virtual Router on page 278
- Enabling PIM for IPv6 on an Interface on page 278
- **ipv6 router pim**

CHAPTER 14

Monitoring PIM for IPv6 Multicast

You can display information about PIM events and parameters. You can use the debug PIM commands to view information about PIM events and the **show ipv6 pim** commands to display information about PIM settings.

Enabling the Display of a PIM Event for IPv6

To display information about the selected event.

- Issue the **debug ipv6 pim** command in Privileged Exec mode:

  host1#**debug ipv6 pim events severity 1 verbosity low**

  Use the **no** version to disable the display.

NOTE: To control the type of events displayed, specify a severity level, and to control how much information to display, specify a verbosity level.

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273

- **debug ipv6 pim**

Disabling the Display of a PIM Event for IPv6

To turn off the display of information previously enabled with the **debug ipv6 pim** command.

- Issue the **undebug ipv6 pim** command in the Privileged Exec mode:

  host1#**undebug ipv6 pim events**

  There is no **no** version to disable the display.

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273
- **undebug ipv6 pim**

Monitoring PIM Router-Level Information for IPv6

Purpose Display general PIM router-level information.

Action To display general PIM router-level information:

```
host1:1#show ipv6 pim
Default PIM Version: 2
Default Domain Id: 0
Default Hello Period: 30
Default Hello HoldTime: 105
Join-Prune Interval: 100
Join-Prune Holdtime: 350
Keepalive Period: 210
Assert Time: 210
Register Suppression Time: 60
Register Probe Time: 5
Register TTL: 64
SSM enabled, range default
Sparse-Mode Graceful Restart Duration: 30
Graceful restart is complete (timer 0 seconds)
Join filter, access-list bronze
Designated Router Priority: 2
```

Meaning Table 79 on page 294 lists the **show ipv6 pim** command output fields.

Table 79: show ipv6 pim

Field Name	Field Description
Default PIM Version	Default PIM version number (always 2)
Default Domain Id	Default Domain Id (always 0)
Default Hello period	Default interval (in minutes) at which the router sends hello messages to neighbors
Default Hello Hold Time	Default time (in minutes) for which the router keeps the neighbor state alive

Table 79: show ipv6 pim *(continued)*

Field Name	Field Description
Join-Prune Interval	Interval value (in seconds) set in the join/prune message originated by the PIM router
Join-Prune Holdtime	Hold time value (in seconds) set in the join/prune message originated by the PIM router. The hold time is 3.5 times the PIM join/prune message interval value.
Keepalive Period	Time SG join state is maintained in the absence of SG Join message
Assert Time	Period after last assert before assert state is timed out
Register Suppression Time	Period during which a designated router stops sending registers to the RP
Register Probe Time	Time before register suppression time (RST) expires when a designated router might send a NULL-Register to the RP
Register TTL	TTL value (in PIM register packets) originated by this PIM router
SSM	State of SSM on this PIM router (enabled or disabled)
range	Default SSM group range or name of the access list specifying the range
Join filter	Name of the join filter access-list (if configured) for this PIM router
Designated Router Priority	Designated router priority value

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273
- **show ipv6 pim**

Monitoring BSR Information for IPv6

Purpose Display BSR information and the group prefixes for which the local router is a candidate RP in a PIM sparse mode environment.

Action To display information on a router that is the elected BSR:

```
host1:1#show ipv6 pim bsr
This PIM router is a Candidate BSR.
  Configured on intf ATM3/0.101, address: ::107:9
  hashMaskLen 30, priority 2, period 60 seconds.
Elected BSR is this router, next BSM in 3 seconds.
```

```
Local candidate RP mapping(s):
Candidate RP ::107:9
::108:86, BSR, hold-time 150, interval 60, priority 192
::108:87, BSR, hold-time 150, interval 60, priority 192, from access-list acl
::108:88, BSR, hold-time 150, interval 60, priority 192, from access-list acl
```

To display information on a router that is a candidate BSR:

```
host1:1#show ipv6 pim bsr
This PIM router is a Candidate BSR.
  Configured on intf ATM3/0.100, address: ::107:9
  hashMaskLen 30, priority 2, period 60 seconds.
Elected BSR is ::107:8 (priority 0), expires in 73 seconds.
```

To display information on a router that is not a candidate BSR:

```
host1:1#show ipv6 pim bsr
This PIM router is not a Candidate BSR.
Elected BSR is ::107:9 (priority 0), expires in 73 seconds.
```

Meaning Table 80 on page 296 lists the **show ipv6 pim bsr** command output fields.

Table 80: show ipv6 pim bsr Output Fields

Field Name	Field Description
Candidacy	Whether or not the router is a candidate BSR
Configured on	Interface on which the router is configured
address	Address of the router
hashMaskLen	Hash mask length
priority	Priority of the router
period	Time between bootstrap messages
Elected BSR	“ this router” or IP address of the elected bootstrap router
next BSM	If BSR is “ this router,” time until the next bootstrap message is sent
expires in	If BSR is not “ this router,” time until the elected BSR expires if no bootstrap messages are received
Local candidate RP mapping(s)	Routers that the mapping agent is evaluating to determine an RP router for this interface

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273
- **show ipv6 pim bsr**

Monitoring PIM Interfaces for IPv6

Purpose Display information about PIM interfaces. Specify no keywords or variables to view information about all PIM interfaces. You can specify the **detail** keyword to view detailed information for all PIM interfaces or for a specified PIM interface. You can specify the **summary** keyword to view the number of configured, enabled, and disabled PIM sparse-mode interfaces. You can specify the **count** keyword to view the number of multicast packets that the interface has sent and received.

NOTE: Provide an interface type and specifier (such as atm 3/0) to display information about that interface only. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.

Action To display information for all PIM interfaces, or for a specified PIM interface:

```
host1# show ipv6 pim interface

PIM Interface Table
Interface Addr Interface   State Ver Mode    Nbr    Hello  J/P    DR    DR Addr
  JoinFilter
                                             Count  Intvl  Intvl  Pri
101::1         ATM2/0.100  Up      2  Sparse  1      30    150     5    101::2
  -
102::1         ATM2/0.101  Up      2  Sparse  1      30    100     2    102::2
  silver
103::1         ATM3/0.102  Up      2  Sparse  1      30    100     2    103::1
  gold
```

To display the number of configured, enabled, and disabled PIM sparse-mode interface:

```
host1#show ipv6 pim interface summary
PIM Interface Summary
  SM:    0, 0 enabled, 0 disabled
```

To display the number of multicast packets that the interface has sent and received:

```
host1#show ipv6 pim interface count
PIM Interface Count
Interface Addr  Interface Name       ControlPktCount In|Out
                                     Hello     JoinPrune  Assert
::107:84        ATM3/0.20             0         0           0
                                      0         0           0
```

Meaning Table 81 on page 297 lists the **show ipv6 pim interface** command output fields.

Table 81: show ipv6 pim interface Output Fields

Field Name	Field Description
Interface Addr	IPv6 address of the interface

Table 81: show ipv6 pim interface Output Fields *(continued)*

Field Name	Field Description
Interface Name	Type and identifier of the interface. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*
Ver	Version of PIM running on this interface
Mode	PIM mode running on this interface: Sparse
Nbr Count	Number of neighbors connected to this interface
Hello Intvl	Time interval at which the interface sends hello messages to neighbors
J/P Interval	Total number of join/prune message interval, in seconds, at which the interface sends the join/prune messages
DR Address	Address of the designated router
SM	Number of PIM sparse mode interfaces: • enabled—Number of interfaces administratively enabled • disabled—Number of interfaces administratively disabled
DR Pri	Designated router priority value
ControlPkt Count In | Out	PIM messages received on and sent from this interface: • Hello—Number of hello messages • JoinPrune—Total number of join and prune messages • Assert—Number of assert messages

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273
- **show ipv6 pim interface**

Monitoring PIM Neighbors for IPv6

Purpose Display information about PIM neighbors that the router discovered. You can specify the **detail** keyword to view detailed information for all PIM neighbors or for a specified PIM neighbor.

NOTE: Provide an interface type and specifier (such as atm 3/0) to display information about that interface only. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.

Action To display information about PIM neighbors that the router discovered:

```
host1#show ipv6 pim neighbor
PIM Neighbor Table
Neighbor Addr   Interface Name        Uptime         Expires  Ver Mode
::107:48        atm2/1.109            1d15:47:35     00:01:41 2   Sparse
::108:58        atm2/1.108            1d15:47:34     00:01:42 2   Sparse
::111:98        atm2/0.110            1d15:48:02     00:01:44 2   Sparse
```

Meaning Table 82 on page 299 lists the **show ipv6 pim neighbor** command output fields.

Table 82: show ipv6 pim neighbor Output Fields

Field Name	Field Description
Neighbor Addr	IPv6 address of the neighbor
Interface Name	Type and specifier of the interface to which the neighbor connects. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.
Uptime	Time since the router discovered this neighbor
Expires	Time available for the neighbor to send a hello message to the interface. If the neighbor does not send a hello message during this time, it will no longer be a neighbor
Ver	Version of PIM that the neighbor is running
Mode	PIM mode that the neighbor is using: sparse

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273
- **show ipv6 pim neighbor**

Monitoring PIM Remote Neighbors for IPv6

Purpose Display information about PIM remote neighbors.

Action To display information about PIM remote neighbors:

```
host1:boston#show ipv6 pim remote-neighbor
PIM RemoteNbr Table
RemoteNbr Addr   OurEnd Addr     Ver Mode        Nbr   Hello  DR Addr
                                                 Count Intvl
1001::1          2001::1         2   Sparse      1     30     ::107:84
      In interface : atm2/1.109
      Out interface: atm2/1.108
```

Meaning Table 83 on page 300 lists the **show ipv6 pim remote-neighbor** command output fields.

Table 83: show ipv6 pim remote-neighbor Output Fields

Field Name	Field Description
Remote Nbr Addr	IPv6 address of remote neighbor
OurEnd Addr	IPv6 address of local interface, such as the local endpoint of a tunnel, that transmits data to remote neighbor
Ver	Version of PIM running on the local interface
Mode	PIM mode running on the local interface; always PIM sparse mode
Nbr Count	Number of remote neighbors detected: 0 or 1
Hello Intvl	Time interval at which the interface sends hello messages to neighbors
DR Addr	Address of designated router
In interface	Type and identifier of the interface on which PIM router receives packets from remote neighbor. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.
Out interface	Type and identifier of the interface on which PIM router sends packets to remote neighbor. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273
- **show ipv6 pim remote-neighbor**

Monitoring PIM Group-to-RP Mappings for IPv6

Purpose Display information about PIM group-to-RP mappings. You can specify the address of a group to view PIM group-to-RP mappings for a particular group. You can specify the **mapping** keyword to display all group-to-RP mappings that the router has recorded.

Action To display information about PIM group-to-RP mappings:

```
host1:8#show ipv6 pim rp mapping
PIM Group-to-RP mapping(s)
Group(s) ff00::/12
  RP ::122:1,  priority 0,  via  static
Group(s) ff0e::1:0/96
  RP ::120:1,  priority 0,  via  static
```

Meaning Table 84 on page 301 lists the **show ipv6 pim rp mapping** command output fields.

Table 84: show ipv6 pim rp mapping Output Fields

Field Name	Field Description
Group	Prefix of the multicast group
RP	IP address of RP router for the multicast group
priority	This field is not functional
via	Method by which the RP router was assigned (static, BSR)

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273
- **show ipv6 pim rp**

Monitoring the RP Router that a Multicast Group Uses for IPv6

Purpose Display the RP Router that a multicast group uses.

Action To display the RP Router that a multicast group uses:

```
host1:2#show ipv6 pim rp-hash 232.1.1.1
Group(s) ff00::/12
  RP ::122:1,  priority 0,  via  static
```

Meaning Table 85 on page 301 lists the **show ipv6 pim rp-hash** command output fields.

Table 85: show ipv6 pim rp-hash Output Fields

Field Name	Field Description
Group	Multicast group
RP	RP router for the multicast group
priority	This field is not functional
via	Method by which the RP router was assigned (static, BSR)

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273
- **show ipv6 pim rp-hash**

Monitoring each (S,G) Pair for PIM Sparse Mode and PIM SSM for IPv6

Purpose Display information for each (S,G) entry for PIM sparse mode and PIM SSM.

Action To display information for each (S,G) entry for PIM sparse mode and PIM SSM:

```
host1:2#show ipv6 pim sparse-mode sg-state
PIM SM route table and oif information
<*, ff0e::1:3>
  Group-to-RP mapping: ff00::/12 RP: ::123:1
  RPF Route: ::123:1/96 IIF: :106:73 UpNbr: ::106:37
  Oifs:
   Address: ::78:7:7 Interface: loopback7
    Local group membership present.
<*, ff0e::a:1>
  Group-to-RP mapping: ff001:/12 RP: ::123:1
  RPF Route: :123:1/96 IIF: :106:73 UpNbr: :106:37
  Oifs:
   Address: ::78:7:7 Interface: loopback7
    Local group membership present.
<::118:34, ff3e::1>
  SSM Group
  RPF Route: ::118:0/96 IIF: :118:1 (Directly attached)
  Oifs:
   Register Oif to RP: ::141:2 suppressed for SSM Group.
   Address: ::134:1 Interface: ATM3/0.104
    Joined as <S, G>   Join Expires: 161
<::118:35, ff3e::1>
  SSM Group
  RPF Route: ::118:0/96 IIF: ::118:1 (Directly attached)
  Oifs:
   Register Oif to RP: ::141:2 suppressed for SSM Group.
   Address: ::134:1 Interface: ATM3/0.104
    Joined as <S, G>   Join Expires: 161
<::10:8, ff0e::5:1>      EntryExpires: 143
  Group-to-RP mapping: ff00::/12 RP: ::123:1
  RPF Route: ::10:0/96 IIF: ::106:73 UpNbr: ::106:37
  Oifs:
   Address: ::78:7:7 Interface: loopback7
    Joined as <*, G>
Count of entries - <S, G>    : 3
                   <*, G>    : 2
                   <*, *, RP>: 0
```

Meaning Table 86 on page 302 lists the **show ipv6 pim sparse-mode sg-state** command output fields.

Table 86: show ipv6 pim sparse-mode sg-state Output Fields

Field Name	Field Description
Group-to-RP mapping	IPv6 addresses and network mask of the multicast group

Table 86: show ipv6 pim sparse-mode sg-state Output Fields *(continued)*

Field Name	Field Description
RP	IPv6 address of RP router
SSM group	Indicates that this is an SSM group
RPF route	IPv6 address and network mask of the RPF route
IIF	IPv6 address of the incoming interface for the RPF route
UpNbr	IPv6 address of the upstream neighbor
Oifs	Outgoing interface
Register Oif to RP	IP address of RP router for the outgoing interface; suppressed for SSM
Address	IPv6 address of outgoing interface
Interface	Type and specifier of the interface. For details about interface types and specifiers, see Interface Types and Specifiers in *JunosE Command Reference Guide*.
Joined as	Type of mapping: • (S,G)—Mapping from a specific source to a specific group • (*,G)—Mapping from any source to a specific group • (*,*,RP)—Mapping from any source to any group
Join expires	Number of seconds before the (S,G) membership expires
Count of entries	Total count of (S,G) pair mappings

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273
- **show ipv6 pim sparse-mode sg-state**

Monitoring Unicast Routes that PIM Sparse Mode Uses for IPv6

Purpose Display the unicast routes that PIM sparse mode is using.

Action To display the unicast routes that PIM sparse mode is using:

```
host1:2#show ipv6 pim sparse-mode unicast-route
PIM SM unicast route table information
Route                        RpfNbr         Iif             Pref  Metric
---------------------------------------------------------------------------
```

```
::122:0      /96                          ::122:1        255   1
Count of entries: 1
```

Meaning Table 87 on page 304 lists the **show ipv6 pim sparse-mode unicast-route** command output fields.

Table 87: show ipv6 pim sparse-mode unicast-route Output Fields

Field Name	Field Description
Route	IPv6 address and network mask for the unicast route
RpfNbr	RPF neighbor
Iif	Incoming interface for the unicast route
Pref	Preference for the unicast route
Metric	Value of metric for the unicast route (type of metric varies with the unicast protocol)
Count of entries	Number of unicast routes that PIM sparse mode is using

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273
- **show ipv6 pim sparse-mode unicast-route**

Monitoring the Threshold for Switching to the Shortest Path Tree at a PIM Designated Router for IPv6

Purpose Display the threshold for switching to the shortest path tree at a PIM designated router.

Action To display the threshold for switching to the shortest path tree at a PIM designated router:

```
host1:2#show ipv6 pim spt-threshold
Access List Name                  SptThreshold(in kbps)
-------------------------------------------------------
1                                 infinity
```

Meaning Table 88 on page 304 lists the **show ipv6 pim spt-threshold** command output fields.

Table 88: show ipv6 pim spt-threshold Output Fields

Field Name	Field Description
Access List Name	Name of the IPv6 access list that specifies the groups to which the threshold applies
SptThreshold (in kbps)	Value at which PIM sparse mode should switch from a shared tree to an SPT. A value of infinity indicates that PIM sparse mode should never switch to an SPT

Related Documentation

- Configuring PIM for IPv6 Multicast on page 273
- **show ipv6 pim spt-threshold**

PART 3

Index

Index

J

L

M