



Junos[®] OS for EX Series Ethernet Switches

Multicast Feature Guide for EX Series Switches

Release

15.1



Modified: 2015-06-12

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Junos[®] OS for EX Series Ethernet Switches Multicast Feature Guide for EX Series Switches
Release 15.1
Copyright © 2015, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

| | | |
|------------------|--|-----------|
| | About the Documentation | xiii |
| | Documentation and Release Notes | xiii |
| | Supported Platforms | xiii |
| | Using the Examples in This Manual | xiii |
| | Merging a Full Example | xiv |
| | Merging a Snippet | xiv |
| | Documentation Conventions | xv |
| | Documentation Feedback | xvii |
| | Requesting Technical Support | xvii |
| | Self-Help Online Tools and Resources | xvii |
| | Opening a Case with JTAC | xviii |
| Part 1 | Overview | |
| Chapter 1 | IGMP Snooping Overview | 3 |
| | IGMP Snooping on EX Series Switches Overview | 3 |
| | How IGMP Snooping Works | 3 |
| | IGMP Message Types | 4 |
| | How Hosts Join and Leave Multicast Groups | 5 |
| | Support for IGMPv3 Multicast Sources | 5 |
| | IGMP Snooping and Forwarding Interfaces | 6 |
| | General Forwarding Rules | 6 |
| | Examples of IGMP Snooping Multicast Forwarding | 7 |
| | Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts | 7 |
| | Scenario 2: Switch Forwarding Multicast Traffic to Another Switch | 8 |
| | Scenario 3: Switch Connected to Hosts Only (No IGMP Querier) | 9 |
| | Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs | 10 |
| Chapter 2 | MLD Snooping Overview | 13 |
| | Understanding MLD Snooping | 13 |
| | How MLD Snooping Works | 13 |
| | MLD Message Types | 14 |
| | How Hosts Join and Leave Multicast Groups | 15 |
| | Support for MLDv2 Multicast Sources | 15 |
| | MLD Snooping and Forwarding Interfaces | 16 |
| | General Forwarding Rules | 17 |

| | | |
|------------------|--|-----------|
| | Examples of MLD Snooping Multicast Forwarding | 17 |
| | Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts | 17 |
| | Scenario 2: Switch Forwarding Multicast Traffic to Another Switch | 18 |
| | Scenario 3: Switch Connected to Hosts Only (No MLD Querier) | 19 |
| | Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs | 20 |
| Chapter 3 | Multicast VLAN Registration Overview | 23 |
| | Understanding Multicast VLAN Registration | 23 |
| | How MVR Works | 23 |
| | MVR Modes | 24 |
| Part 2 | Configuration | |
| Chapter 4 | Configuration Examples | 27 |
| | Example: Configuring IGMP Snooping on EX Series Switches | 27 |
| | Example: Configuring MLD Snooping | 30 |
| | Example: Configuring Multicast VLAN Registration | 33 |
| Chapter 5 | Configuration Tasks | 39 |
| | Configuring IGMP Snooping (CLI Procedure) | 39 |
| | Enabling or Disabling IGMP Snooping on VLANs | 41 |
| | Configuring the IGMP Version | 41 |
| | Enabling Immediate Leave | 42 |
| | Configuring an Interface as a Multicast-Router Interface | 43 |
| | Configuring Static Group Membership on an Interface | 44 |
| | Changing the Timer and Counter Values | 45 |
| | Configuring IGMP Snooping (J-Web Procedure) | 46 |
| | Configuring IGMP Snooping Tracing Operations (CLI Procedure) | 48 |
| | Configuring Tracing Operations | 49 |
| | Viewing, Stopping, and Restarting Tracing Operations | 50 |
| | Configuring MLD Snooping on a VLAN (CLI Procedure) | 50 |
| | Enabling or Disabling MLD Snooping on VLANs | 52 |
| | Configuring the MLD Version | 53 |
| | Enabling Immediate Leave | 54 |
| | Configuring an Interface as a Multicast-Router Interface | 54 |
| | Configuring Static Group Membership on an Interface | 55 |
| | Changing the Timer and Counter Values | 56 |
| | Configuring MLD Snooping Tracing Operations (CLI Procedure) | 58 |
| | Configuring Tracing Operations | 59 |
| | Viewing, Stopping, and Restarting Tracing Operations | 59 |
| | Configuring Multicast VLAN Registration (CLI Procedure) | 60 |
| Chapter 6 | Configuration Statements | 63 |
| | [edit protocols] Configuration Statement Hierarchy on EX Series Switches | 65 |
| | [edit protocols igmp] Configuration Statement Hierarchy on EX Series Switches | 67 |
| | Supported Statements in the [edit protocols igmp] Hierarchy Level | 67 |
| | Unsupported Statements in the [edit protocols igmp] Hierarchy Level | 68 |

| | |
|---|-----|
| [edit protocols igmp-snooping] Configuration Statement Hierarchy on EX Series | |
| Switches | 68 |
| Supported Statements in the [edit protocols igmp-snooping] Hierarchy | |
| Level | 68 |
| Unsupported Statements in the [edit protocols igmp-snooping] Hierarchy | |
| Level | 69 |
| accounting (Protocols IGMP Interface) | 70 |
| accounting (Protocols IGMP) | 70 |
| address (Anycast RPs) | 71 |
| address (Local RPs) | 72 |
| anycast-pim | 73 |
| assert-timeout | 74 |
| auto-rp | 75 |
| bootstrap | 76 |
| bootstrap-export | 77 |
| bootstrap-import | 78 |
| bootstrap-priority | 79 |
| data-forwarding | 80 |
| dense-groups | 81 |
| disable (IGMP Snooping) | 81 |
| disable (Protocols IGMP) | 82 |
| disable (MLD Snooping) | 82 |
| disable (PIM) | 83 |
| dr-election-on-p2p | 84 |
| dr-register-policy | 84 |
| embedded-rp | 85 |
| export (Protocols PIM Bootstrap) | 86 |
| family (Bootstrap) | 87 |
| family (Local RP) | 88 |
| graceful-restart (Protocols PIM) | 89 |
| group (IGMP Snooping) | 89 |
| group (Protocols IGMP) | 90 |
| group (MLD Snooping) | 91 |
| group-ranges | 92 |
| groups (Multicast VLAN Registration) | 93 |
| hello-interval (Protocols PIM) | 94 |
| hold-time (Protocols PIM) | 95 |
| igmp-snooping | 96 |
| immediate-leave (Protocols IGMP) | 97 |
| immediate-leave (IGMP Snooping) | 98 |
| immediate-leave (MLD Snooping) | 99 |
| import (Protocols PIM Bootstrap) | 100 |
| import (Protocols PIM) | 101 |
| infinity | 102 |
| install (Multicast VLAN Registration) | 102 |
| interface (IGMP Snooping) | 103 |
| interface (Protocols PIM) | 104 |
| interface (Protocols IGMP) | 106 |
| interface (MLD Snooping) | 107 |

| | |
|---|-----|
| join-load-balance | 108 |
| local | 109 |
| local-address (Protocols PIM) | 110 |
| mapping-agent-election | 111 |
| maximum-rps | 112 |
| mld-snooping | 113 |
| mode (Protocols PIM) | 114 |
| multicast-router-interface (IGMP Snooping) | 115 |
| multicast-router-interface (MLD Snooping) | 116 |
| neighbor-policy | 117 |
| pim | 118 |
| priority (PIM Interfaces) | 122 |
| priority (Bootstrap) | 123 |
| priority (PIM RPs) | 124 |
| promiscuous-mode (Protocols IGMP) | 125 |
| proxy (Multicast VLAN Registration) | 125 |
| query-interval (Protocols IGMP) | 126 |
| query-last-member-interval (Protocols IGMP) | 127 |
| query-response-interval (Protocols IGMP) | 128 |
| receiver | 129 |
| restart-duration (Protocols PIM) | 130 |
| rib-group (Protocols PIM) | 131 |
| robust-count (IGMP Snooping) | 132 |
| robust-count (Protocols IGMP) | 132 |
| robust-count (MLD Snooping) | 133 |
| rp | 134 |
| rp-register-policy | 136 |
| rp-set | 137 |
| source (Multicast VLAN Registration) | 137 |
| source (Protocols IGMP) | 138 |
| source-vlans | 139 |
| spt-threshold | 140 |
| ssm-map (Protocols IGMP) | 141 |
| static (IGMP Snooping) | 141 |
| static (Protocols PIM) | 142 |
| static (Protocols IGMP) | 143 |
| static (MLD Snooping) | 144 |
| traceoptions (Protocols PIM) | 145 |
| traceoptions (Protocols IGMP) | 148 |
| traceoptions (IGMP Snooping) | 151 |
| traceoptions (MLD Snooping) | 153 |
| version (Protocols IGMP) | 155 |
| version (IGMP Snooping) | 156 |
| version (MLD Snooping) | 157 |
| version (PIM) | 158 |
| vlan (IGMP Snooping) | 159 |
| vlan (MLD Snooping) | 161 |

| | | |
|------------------|---|------------|
| Part 3 | Administration | |
| Chapter 7 | Routine Monitoring | 165 |
| | Monitoring IGMP Snooping | 165 |
| | Verifying IGMP Snooping (CLI Procedure) | 168 |
| | Verifying IGMP Snooping Memberships | 168 |
| | Viewing IGMP Snooping Statistics | 169 |
| | Viewing IGMP Snooping Routing Information | 169 |
| | Verifying MLD Snooping (CLI Procedure) | 170 |
| | Verifying MLD Snooping Memberships | 170 |
| | Verifying MLD Snooping VLANs | 171 |
| | Viewing MLD Snooping Statistics | 172 |
| | Viewing MLD Snooping Routing Information | 172 |
| Chapter 8 | Operational Commands | 175 |
| | clear igmp membership | 177 |
| | clear igmp statistics | 180 |
| | clear igmp-snooping membership | 182 |
| | clear igmp-snooping statistics | 183 |
| | clear mld-snooping membership | 184 |
| | clear mld-snooping statistics | 185 |
| | clear multicast bandwidth-admission | 186 |
| | clear multicast scope | 188 |
| | clear multicast sessions | 189 |
| | clear multicast statistics | 190 |
| | clear pim join | 191 |
| | clear pim register | 193 |
| | clear pim statistics | 195 |
| | mtrace | 198 |
| | mtrace from-source | 201 |
| | mtrace monitor | 204 |
| | mtrace to-gateway | 206 |
| | show igmp group | 209 |
| | show igmp interface | 213 |
| | show igmp statistics | 217 |
| | show igmp-snooping membership | 220 |
| | show igmp-snooping route | 223 |
| | show igmp-snooping statistics | 225 |
| | show igmp-snooping vlans | 227 |
| | show mld-snooping membership | 230 |
| | show mld-snooping route | 233 |
| | show mld-snooping statistics | 236 |
| | show mld-snooping vlans | 238 |
| | show multicast flow-map | 241 |
| | show multicast interface | 243 |
| | show multicast mrinfo | 245 |
| | show multicast next-hops | 247 |
| | show multicast pim-to-igmp-proxy | 250 |
| | show multicast pim-to-mld-proxy | 252 |
| | show multicast route | 254 |

| | |
|-------------------------------|-----|
| show multicast rpf | 261 |
| show multicast scope | 265 |
| show multicast sessions | 267 |
| show multicast usage | 270 |
| show pim bootstrap | 273 |
| show pim interfaces | 275 |
| show pim join | 278 |
| show pim neighbors | 300 |
| show pim rps | 304 |
| show pim source | 311 |
| show pim statistics | 314 |

List of Figures

| | | |
|------------------|---|-----------|
| Part 1 | Overview | |
| Chapter 1 | IGMP Snooping Overview | 3 |
| | Figure 1: Multicast Traffic Flow with IGMP Snooping Enabled | 4 |
| | Figure 2: Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts | 8 |
| | Figure 3: Scenario 2: Switch Forwarding Multicast Traffic to Another Switch | 9 |
| | Figure 4: Scenario 3: Switch Connected to Hosts Only (No IGMP Querier) | 10 |
| | Figure 5: Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs | 11 |
| Chapter 2 | MLD Snooping Overview | 13 |
| | Figure 6: Multicast Traffic Flow with MLD Snooping Enabled | 14 |
| | Figure 7: Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts | 18 |
| | Figure 8: Scenario 2: Switch Forwarding Multicast Traffic to Another Switch | 19 |
| | Figure 9: Scenario 3: Switch Connected to Hosts Only (No MLD Querier) | 20 |
| | Figure 10: Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs | 21 |
| Part 2 | Configuration | |
| Chapter 4 | Configuration Examples | 27 |
| | Figure 11: Example IGMP Snooping Topology | 28 |
| | Figure 12: Example MLD Snooping Topology | 31 |
| | Figure 13: MVR Topology in Transparent Mode | 35 |
| | Figure 14: MVR Topology in Proxy Mode | 36 |

List of Tables

| | | |
|------------------|--|-------------|
| | About the Documentation | xiii |
| | Table 1: Notice Icons | xv |
| | Table 2: Text and Syntax Conventions | xv |
| Part 2 | Configuration | |
| Chapter 5 | Configuration Tasks | 39 |
| | Table 3: IGMP Snooping Configuration Fields | 47 |
| | Table 4: Supported Tracing Operations for IGMP Snooping | 48 |
| | Table 5: Supported Tracing Operations for MLD Snooping | 58 |
| Part 3 | Administration | |
| Chapter 7 | Routine Monitoring | 165 |
| | Table 6: Summary of IGMP Snooping Output Fields | 166 |
| Chapter 8 | Operational Commands | 175 |
| | Table 7: mtrace Output Fields | 198 |
| | Table 8: mtrace from-source Output Fields | 202 |
| | Table 9: mtrace monitor Output Fields | 204 |
| | Table 10: mtrace to-gateway Output Fields | 207 |
| | Table 11: show igmp group Output Fields | 209 |
| | Table 12: show igmp interface Output Fields | 213 |
| | Table 13: show igmp statistics Output Fields | 217 |
| | Table 14: show igmp-snooping membership Output Fields | 220 |
| | Table 15: show igmp-snooping route Output Fields | 223 |
| | Table 16: show igmp-snooping statistics Output Fields | 225 |
| | Table 17: show igmp-snooping vlans Output Fields | 227 |
| | Table 18: show mld-snooping membership Output Fields | 230 |
| | Table 19: show mld-snooping route Output Fields | 233 |
| | Table 20: show mld-snooping statistics Output Fields | 236 |
| | Table 21: show mld-snooping vlans Output Fields | 238 |
| | Table 22: show multicast flow-map Output Fields | 241 |
| | Table 23: show multicast interface Output Fields | 243 |
| | Table 24: show multicast minfo Output Fields | 245 |
| | Table 25: show multicast next-hops Output Fields | 248 |
| | Table 26: show multicast pim-to-igmp-proxy Output Fields | 250 |
| | Table 27: show multicast pim-to-mld-proxy Output Fields | 252 |
| | Table 28: show multicast route Output Fields | 255 |
| | Table 29: show multicast rpf Output Fields | 262 |
| | Table 30: show multicast scope Output Fields | 265 |

| | |
|---|-----|
| Table 31: show multicast sessions Output Fields | 267 |
| Table 32: show multicast usage Output Fields | 271 |
| Table 33: show pim bootstrap Output Fields | 273 |
| Table 34: show pim interfaces Output Fields | 275 |
| Table 35: show pim join Output Fields | 280 |
| Table 36: show pim neighbors Output Fields | 301 |
| Table 37: show pim rps Output Fields | 305 |
| Table 38: show pim source Output Fields | 312 |
| Table 39: show pim statistics Output Fields | 315 |

About the Documentation

- Documentation and Release Notes on page xiii
- Supported Platforms on page xiii
- Using the Examples in This Manual on page xiii
- Documentation Conventions on page xv
- Documentation Feedback on page xvii
- Requesting Technical Support on page xvii

Documentation and Release Notes

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

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

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

Supported Platforms

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

- EX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

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

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

| Convention | Description | Examples |
|--------------------------------|---|--|
| Fixed-width text like this | Represents output that appears on the terminal screen. | <code>user@host> show chassis alarms</code> <code>No alarms currently active</code> |
| <i>Italic text like this</i> | <ul style="list-style-type: none">Introduces or emphasizes important new terms.Identifies guide names.Identifies RFC and Internet draft titles. | <ul style="list-style-type: none">A policy <i>term</i> is a named structure that defines match conditions and actions.<i>Junos OS CLI User Guide</i>RFC 1997, <i>BGP Communities Attribute</i> |
| <i>Italic text like this</i> | Represents variables (options for which you substitute a value) in commands or configuration statements. | Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i> |
| Text like this | Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components. | <ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE. |
| < > (angle brackets) | Encloses optional keywords or variables. | stub <default-metric <i>metric</i>>; |
| (pipe symbol) | Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity. | broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>) |
| # (pound sign) | Indicates a comment specified on the same line as the configuration statement to which it applies. | rsvp { # Required for dynamic MPLS only |
| [] (square brackets) | Encloses a variable for which you can substitute one or more values. | community name members [<i>community-ids</i>] |
| Indentation and braces ({ }) | Identifies a level in the configuration hierarchy. | [edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } } |
| ;(semicolon) | Identifies a leaf statement at a configuration hierarchy level. | } |
| GUI Conventions | | |
| Bold text like this | Represents graphical user interface (GUI) items you click or select. | <ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel. |

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

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

Documentation Feedback

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

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [IGMP Snooping Overview on page 3](#)
- [MLD Snooping Overview on page 13](#)
- [Multicast VLAN Registration Overview on page 23](#)

CHAPTER 1

IGMP Snooping Overview

- [IGMP Snooping on EX Series Switches Overview on page 3](#)

IGMP Snooping on EX Series Switches Overview

Internet Group Management Protocol (IGMP) snooping constrains the flooding of IPv4 multicast traffic on VLANs on a switch. When IGMP snooping is enabled on a VLAN, a Juniper Networks EX Series Ethernet Switch examines IGMP messages between hosts and multicast routers and learns which hosts are interested in receiving traffic for a multicast group. Based on what it learns, the switch then forwards multicast traffic only to those interfaces in the VLAN that are connected to interested receivers instead of flooding the traffic to all interfaces.

IGMP snooping on EX Series switches supports IGMP version 1 (IGMPv1), IGMPv2, and IGMPv3. For details on IGMP, see the following standards:

- IGMPv1—See RFC 1112, *Host extensions for IP multicasting*.
- IGMPv2—See RFC 2236, *Internet Group Management Protocol, Version 2*.
- For IGMPv3—See RFC 3376, *Internet Group Management Protocol, Version 3*.

This topic covers:

- [How IGMP Snooping Works on page 3](#)
- [IGMP Message Types on page 4](#)
- [How Hosts Join and Leave Multicast Groups on page 5](#)
- [Support for IGMPv3 Multicast Sources on page 5](#)
- [IGMP Snooping and Forwarding Interfaces on page 6](#)
- [General Forwarding Rules on page 6](#)
- [Examples of IGMP Snooping Multicast Forwarding on page 7](#)


How IGMP Snooping Works

A Layer 2 switch usually learns *unicast* media access control (MAC) addresses by checking the source address field of the frames it receives. However, a *multicast* MAC address can never be the source address for a packet. As a result, the switch floods multicast traffic on the VLAN, consuming significant amounts of bandwidth.

You can enable IGMP snooping on a switch to avoid this flooding. When IGMP snooping is enabled, the switch monitors IGMP messages between receivers and multicast routers and uses the content of the messages to build an IPv4 multicast forwarding table—a database of multicast groups and the interfaces that are connected to members of the groups. When the switch receives multicast traffic for a multicast group, it uses the forwarding table to forward the traffic only to interfaces that are connected to receivers that belong to the multicast group.

Figure 1 on page 4 shows an example of multicast traffic flow with IGMP snooping enabled.

Figure 1: Multicast Traffic Flow with IGMP Snooping Enabled



IGMP Message Types

Multicast routers use IGMP to learn, for each of their attached physical networks, which groups have interested listeners. In any given subnet, one multicast router acts as an IGMP querier. The IGMP querier sends out the following types of queries to hosts:

- General query—Asks whether any host is listening to any group.
- Group-specific query—(IGMPv2 and IGMPv3 only) Asks whether any host is listening to a specific multicast group. This query is sent in response to a host leaving the multicast group and allows the router to quickly determine if any remaining hosts are interested in the group.
- Group-and-source-specific query—(IGMPv3 only) Asks whether any host is listening to group multicast traffic from a specific multicast source. This query is sent in response to a host indicating that it is no longer interested in receiving group multicast traffic from the multicast source and allows the router to quickly determine any remaining hosts are interested in receiving group multicast traffic from that source.

Hosts that are multicast listeners send the following kinds of messages:

- Membership report—Indicates that the host wants to join a particular multicast group.
- Leave report—(IGMPv2 and IGMPv3 only) Indicates that the host wants to leave a particular multicast group.

How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in either of two ways:

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

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

Hosts can leave a multicast group in either of two ways:

- By not responding to periodic queries within a set interval of time. This results in what is known as a “silent leave.” This is the only method available to IGMPv1 hosts.
- By sending a leave report. This method can be used by IGMPv2 and IGMPv3 hosts.



NOTE: If a host is connected to the switch through a hub, the host does not automatically leave the multicast group if it disconnects from the hub. The host remains a member of the group until group membership times out and a silent leave occurs. If another host connects to the hub port before the silent leave occurs, the new host might receive the group multicast traffic until the silent leave, even though it never sent an membership report.

Support for IGMPv3 Multicast Sources

In IGMPv3, a host can send a membership report that includes a list of source addresses. When the host sends a membership report in INCLUDE mode, the host is interested in group multicast traffic only from those sources in the source address list. If host sends a membership report in EXCLUDE mode, the host is interested in group multicast traffic from any source *except* the sources in the source address list. A host can also send an EXCLUDE report in which the source-list parameter is empty, which is known as an EXCLUDE NULL report. An EXCLUDE NULL report indicates that the host wants to join the multicast group and receive packets from all sources.

EX Series switches support IGMPv3 membership reports that are in INCLUDE and EXCLUDE mode. However, EX Series switches do not support forwarding on a per-source basis. Instead, a switch consolidates all INCLUDE and EXCLUDE mode reports it receives on a VLAN for a specified group into a single route that includes all multicast sources for

that group, with the next hop being all interfaces that have interested receivers for the group. As a result, interested receivers on the VLAN can receive traffic from a source that they did not include in their INCLUDE report or from a source they excluded in their EXCLUDE report. For example, if Host 1 wants traffic for G from Source A and Host 2 wants traffic for G from Source B, they both receive traffic for G regardless of whether A or B sends the traffic.

IGMP Snooping and Forwarding Interfaces

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

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

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

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



NOTE: For a switch to learn multicast-router interfaces and group-member interfaces, an IGMP querier must exist in the network. For the switch itself to function as an IGMP querier, IGMP must be enabled on the switch.

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

General Forwarding Rules

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

IGMP traffic is forwarded as follows:

- IGMP general queries received on a multicast-router interface are forwarded to all other interfaces in the VLAN.
- IGMP group-specific queries received on a multicast-router interface are forwarded to only those interfaces in the VLAN that are members of the group.

- IGMP reports received on a host interface are forwarded to multicast-router interfaces in the same VLAN, but not to the other host interfaces in the VLAN.

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

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

Examples of IGMP Snooping Multicast Forwarding

The following examples are provided to illustrate how IGMP snooping forwards multicast traffic in different topologies:

- [Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts on page 7](#)
- [Scenario 2: Switch Forwarding Multicast Traffic to Another Switch on page 8](#)
- [Scenario 3: Switch Connected to Hosts Only \(No IGMP Querier\) on page 9](#)
- [Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs on page 10](#)

Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts


In the topology shown in [Figure 2 on page 8](#), a switch acting as a Layer 2 device receives multicast traffic belonging to multicast group **239.10.1.1** from Source A, which is connected to the multicast router. It also receives multicast traffic belonging to multicast group **225.100.100.1** from Source B, which is connected directly to the switch. All interfaces on the switch belong to the same VLAN.

Because the switch receives IGMP queries from the multicast router on interface P1, IGMP snooping learns that interface P1 is a multicast-router interface and adds the interface to its multicast cache table. It forwards any IGMP general queries it receives on this interface to all host interfaces on the switch, and, in turn, forwards membership reports it receives from hosts to the multicast-router interface.

In the example, Hosts A and C have responded to the membership queries with membership reports for group **239.10.1.1**. IGMP snooping adds interfaces P2 and P4 to its multicast cache table as member interfaces for group **239.10.1.1**. It forwards the group multicast traffic received from Source A to Hosts A and C, but not to Hosts B and D.

Host B has responded to the membership queries with a membership report for group **225.100.100.1**. The switch adds interface P3 to its multicast cache table as a member interface for group **225.100.100.1** and forwards multicast traffic it receives from Source B to Host B. The switch also forwards the multicast traffic it receives from Source B to the multicast-router interface P1.

Figure 2: Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts




Scenario 2: Switch Forwarding Multicast Traffic to Another Switch

In the topology shown in [Figure 3 on page 9](#), a multicast source is connected to Switch A. Switch A in turn is connected to another switch, Switch B. Hosts on both Switch A and B are potential members of the multicast group. Both switches are acting as Layer 2 devices and all interfaces on the switches are members of the same VLAN.

Switch A receives IGMP queries from the multicast router on interface P1, making interface P1 a multicast-router interface for Switch A. Switch A forwards all general IGMP queries it receives on this interface to the other interfaces on the switch, including the interface connecting Switch B. Because Switch B receives the forwarded IGMP queries on interface P6, P6 is the multicast-router interface for Switch B. Switch B forwards the group membership report it receives from Host C to Switch A through its multicast-router interface. Switch A forwards the membership report to its multicast-router interface, includes interface P5 in its multicast cache table as a group-member interface, and forwards multicast traffic from the source to Switch B.

Figure 3: Scenario 2: Switch Forwarding Multicast Traffic to Another Switch



In certain implementations, you might have to configure P6 on Switch B as a static multicast-router interface to avoid a delay in a host receiving multicast traffic. For example, if Switch B receives unsolicited membership reports from its hosts before it learns which interface is its multicast-router interface, it does not forward those reports to Switch A. If Switch A then receives multicast traffic, it does not forward the traffic to Switch B, because it has not received any membership reports on interface P5. This issue will resolve when the multicast router sends out its next general query; however, it can cause a delay in the host receiving multicast traffic. You can statically configure interface P6 as a multicast-router interface to solve this issue.


Scenario 3: Switch Connected to Hosts Only (No IGMP Querier)

In the topology shown in [Figure 4 on page 10](#), a switch is connected to a multicast source and to hosts. There is no multicast router in this topology—hence there is no IGMP querier. Without an IGMP querier to respond to, a host does not send periodic membership reports. As a result, even if the host sends an unsolicited join to join a multicast group, its membership in the multicast group times out.

For IGMP snooping to work correctly in this network so that the switch forwards multicast traffic to Hosts A and C only, you can either:

- Configure interfaces P2 and P4 as static group-member interfaces.
- Configure a routed VLAN interface (RVI) on the VLAN and enable IGMP on it. In this case, the switch itself acts as an IGMP querier, and the hosts can dynamically join the multicast group and refresh their group membership by responding to the queries.

Figure 4: Scenario 3: Switch Connected to Hosts Only (No IGMP Querier)




Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs

In the topology shown in [Figure 5 on page 11](#), a multicast source, Multicast Router A, and Hosts A and B are connected to the switch and are in VLAN 10. Multicast Router B and Hosts C and D are also connected to the switch and are in VLAN 20.

In a pure Layer 2 environment, traffic is not forwarded between VLANs. For Host C to receive the multicast traffic from the source on VLAN 10, RVIs must be created on VLAN 10 and VLAN 20 to permit routing of the multicast traffic between the VLANs. In addition, PIM must be enabled on the switch to perform the multicast routing.

Figure 5: Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs



Related Documentation

- [Understanding Multicast VLAN Registration on page 23](#)
- [Example: Configuring IGMP Snooping on EX Series Switches on page 27](#)
- [Configuring IGMP Snooping \(CLI Procedure\) on page 39](#)
- [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#)

CHAPTER 2

MLD Snooping Overview

- [Understanding MLD Snooping on page 13](#)

Understanding MLD Snooping

Multicast Listener Discovery (MLD) snooping constrains the flooding of IPv6 multicast traffic on VLANs. When MLD snooping is enabled on a VLAN, a Juniper Networks EX Series Ethernet Switch examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving traffic for a multicast group. On the basis of what it learns, the switch then forwards multicast traffic only to those interfaces in the VLAN that are connected to interested receivers instead of flooding the traffic to all interfaces.

MLD snooping supports MLD version 1 (MLDv1) and MLDv2. For details on MLDv1 and MLDv2, see the following standards:

- MLDv1—See RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*.
- MLDv2—See RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*.

This topic covers:

- [How MLD Snooping Works on page 13](#)
- [MLD Message Types on page 14](#)
- [How Hosts Join and Leave Multicast Groups on page 15](#)
- [Support for MLDv2 Multicast Sources on page 15](#)
- [MLD Snooping and Forwarding Interfaces on page 16](#)
- [General Forwarding Rules on page 17](#)
- [Examples of MLD Snooping Multicast Forwarding on page 17](#)

How MLD Snooping Works


By default, a switch floods Layer 2 multicast traffic on all of the interfaces belonging to that VLAN on a switch, except for the interface that is the source of the multicast traffic. This behavior can consume significant amounts of bandwidth.

You can enable MLD snooping to avoid this flooding. When you enable MLD snooping, the switch monitors MLD messages between receivers (hosts) and multicast routers

and uses the content of the messages to build an IPv6 multicast forwarding table—a database of IPv6 multicast groups and the interfaces that are connected to the interested members of each group. When the switch receives multicast traffic for a multicast group, it uses the forwarding table to forward the traffic only to interfaces that are connected to receivers that belong to the multicast group.

Figure 6 on page 14 shows an example of multicast traffic flow with MLD snooping enabled.

Figure 6: Multicast Traffic Flow with MLD Snooping Enabled



MLD Message Types

Multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. In any given subnet, one multicast router is elected to act as an MLD querier. The MLD querier sends out the following types of queries to hosts:

- General query—Asks whether any host is listening to any group.
- Group-specific query—Asks whether any host is listening to a specific multicast group. This query is sent in response to a host leaving the multicast group and allows the router to quickly determine if any remaining hosts are interested in the group.

- Group-and-source-specific query—(MLD version 2 only) Asks whether any host is listening to group multicast traffic from a specific multicast source. This query is sent in response to a host indicating that it is no longer interested in receiving group multicast traffic from the multicast source and allows the router to quickly determine any remaining hosts are interested in receiving group multicast traffic from that source.

Hosts that are multicast listeners send the following kinds of messages:

- Membership report—Indicates that the host wants to join a particular multicast group.
- Leave report—Indicates that the host wants to leave a particular multicast group.

Strictly speaking, only MLDv1 hosts use two different kinds of reports to indicate whether they want to join or leave a group. MLDv2 hosts send only one kind of report, the contents of which indicate whether they want to join or leave a group. However, for simplicity's sake, the MLD snooping documentation uses the term *membership report* for a report that indicates that a host wants to join a group and uses the term *leave report* for a report that indicates a host wants to leave a group.

How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in either of two ways:

- By sending an unsolicited membership report that specifies the multicast group that the host is attempting to join.
- By sending a membership report in response to a query from a multicast router.

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

Hosts can leave multicast groups in either of two ways:

- By not responding to periodic queries within a set interval of time. This results in what is known as a “silent leave.”
- By sending a leave report.



NOTE: If a host is connected to the switch through a hub, the host does not automatically leave the multicast group if it disconnects from the hub. The host remains a member of the group until group membership times out and a silent leave occurs. If another host connects to the hub port before the silent leave occurs, the new host might receive the group multicast traffic until the silent leave, even though it never sent an membership report.

Support for MLDv2 Multicast Sources

In MLDv2, a host can send a membership report that includes a list of source addresses. When the host sends a membership report in INCLUDE mode, the host is interested in

group multicast traffic only from those sources in the source address list. If host sends a membership report in EXCLUDE mode, the host is interested in group multicast traffic from any source *except* the sources in the source address list. A host can also send an EXCLUDE report in which the source-list parameter is empty, which is known as an EXCLUDE NULL report. An EXCLUDE NULL report indicates that the host wants to join the multicast group and receive packets from all sources.

The switch support MLDv2 membership reports that are in INCLUDE and EXCLUDE mode. However, the switch does not support forwarding on a per-source basis. Instead, a switch consolidates all INCLUDE and EXCLUDE mode reports it receives on a VLAN for a specified group into a single route that includes all multicast sources for that group, with the next hop being all interfaces that have interested receivers for the group. As a result, interested receivers on the VLAN can receive traffic from a source that they did not include in their INCLUDE report or from a source they excluded in their EXCLUDE report. For example, if Host 1 wants traffic for group G from Source A and Host 2 wants traffic for group G from Source B, they both receive traffic for group G regardless of whether A or B sends the traffic.

MLD Snooping and Forwarding Interfaces

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

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

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

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



NOTE: For a switch to learn multicast-router interfaces and group-member interfaces, an MLD querier must exist in the network. For the switch itself to function as an MLD querier, MLD must be enabled on the switch.

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

General Forwarding Rules

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

MLD protocol traffic is forwarded as follows:

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

Multicast traffic that is not MLD protocol traffic is forwarded as follows:

- An unregistered multicast packet—that is, a packet for a group that has no current members—is forwarded to all multicast-router interfaces in the VLAN.
- A registered multicast packet is forwarded only to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN.

Examples of MLD Snooping Multicast Forwarding

The following examples are provided to illustrate how MLD snooping forwards multicast traffic in different topologies:

- [Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts on page 17](#)
- [Scenario 2: Switch Forwarding Multicast Traffic to Another Switch on page 18](#)
- [Scenario 3: Switch Connected to Hosts Only \(No MLD Querier\) on page 19](#)
- [Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs on page 20](#)

Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts

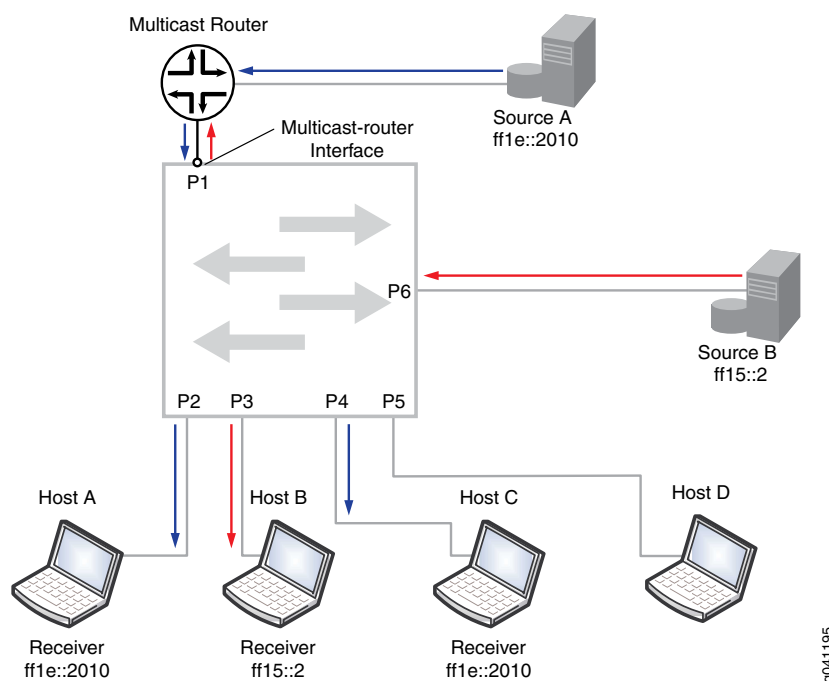
In the topology shown in [Figure 7 on page 18](#), a switch acting as a Layer 2 device receives multicast traffic belonging to multicast group **ff1e::2010** from Source A, which is connected to the multicast router. It also receives multicast traffic belonging to multicast group **ff15::2** from Source B, which is connected directly to the switch. All interfaces on the switch belong to the same VLAN.

Because the switch receives MLD queries from the multicast router on interface P1, MLD snooping learns that interface P1 is a multicast-router interface and adds the interface to its multicast forwarding table. It forwards any MLD general queries it receives on this interface to all host interfaces on the switch, and, in turn, forwards membership reports it receives from hosts to the multicast-router interface.

In the example, Hosts A and C have responded to the general queries with membership reports for group **ff1e::2010**. MLD snooping adds interfaces P2 and P4 to its multicast forwarding table as member interfaces for group **ff1e::2010**. It forwards the group multicast traffic received from Source A to Hosts A and C, but not to Hosts B and D.

Host B has responded to the general queries with a membership report for group **ff15::2**. The switch adds interface P3 to its multicast forwarding table as a member interface for group **ff15::2** and forwards multicast traffic it receives from Source B to Host B. The switch also forwards the multicast traffic it receives from Source B to the multicast-router interface P1.

Figure 7: Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts



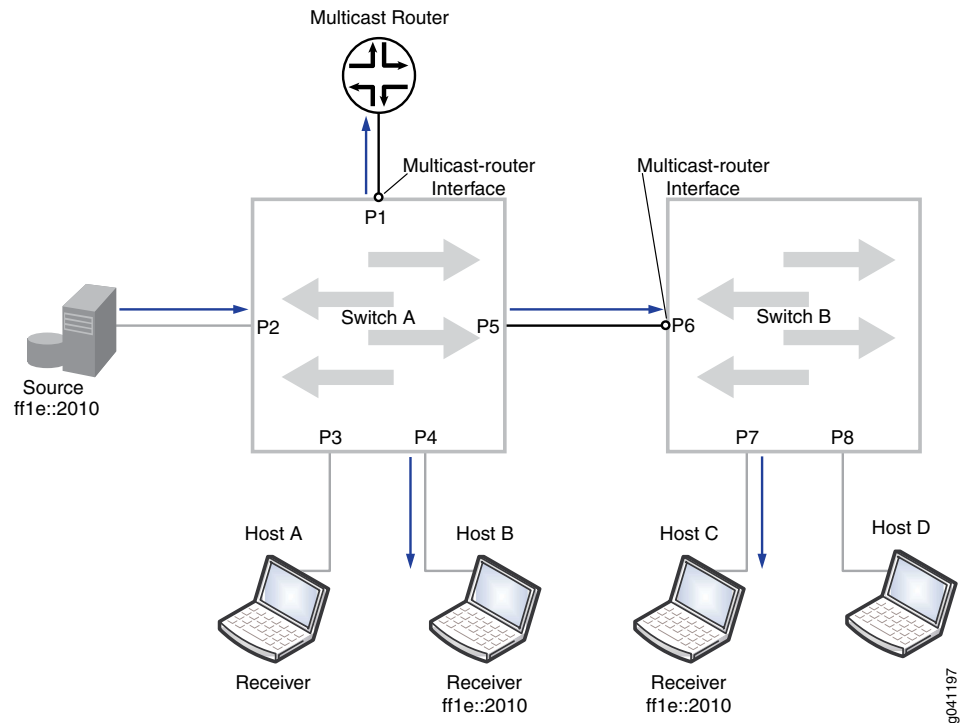
Scenario 2: Switch Forwarding Multicast Traffic to Another Switch

In the topology show in [Figure 8 on page 19](#), a multicast source is connected to Switch A. Switch A in turn is connected to another switch, Switch B. Hosts on both Switch A and B are potential members of the multicast group. Both switches are acting as Layer 2 devices, and all interfaces on the switches are members of the same VLAN.

Switch A receives MLD queries from the multicast router on interface P1, making interface P1 a multicast-router interface for Switch A. Switch A forwards all general queries it receives on this interface to the other interfaces on the switch, including the interface connecting Switch B. Because Switch B receives the forwarded MLD queries on interface P6, P6 is the multicast-router interface for Switch B. Switch B forwards the membership report it receives from Host C to Switch A through its multicast-router interface. Switch A forwards the membership report to its multicast-router interface, includes interface

P5 in its multicast forwarding table as a group-member interface, and forwards multicast traffic from the source to Switch B.

Figure 8: Scenario 2: Switch Forwarding Multicast Traffic to Another Switch



In certain implementations, you might have to configure P6 on Switch B as a static multicast-router interface to avoid a delay in a host receiving multicast traffic. For example, if Switch B receives unsolicited membership reports from its hosts before it learns which interface is its multicast-router interface, it does not forward those reports to Switch A. If Switch A then receives multicast traffic, it does not forward the traffic to Switch B, because it has not received any membership reports on interface P5. This issue will resolve when the multicast router sends out its next general query; however, it can cause a delay in the host receiving multicast traffic. You can statically configure interface P6 as a multicast-router interface to solve this issue.

Scenario 3: Switch Connected to Hosts Only (No MLD Querier)


In the topology shown in [Figure 9 on page 20](#), a switch is connected to a multicast source and to hosts. There is no multicast router in this topology—hence there is no MLD querier. Without an MLD querier to respond to, a host does not send periodic membership reports. As a result, even if the host sends an unsolicited membership report to join a multicast group, its membership in the multicast group will time out.

For MLD snooping to work correctly in this network so that the switch forwards multicast traffic to Hosts A and C only, you can either:

- Configure interfaces P2 and P4 as static group-member interfaces.

- Configure a routed VLAN interface (RVI) on the VLAN and enable MLD on it. In this case, the switch itself acts as an MLD querier, and the hosts can dynamically join the multicast group and refresh their group membership by responding to the queries.

Figure 9: Scenario 3: Switch Connected to Hosts Only (No MLD Querier)

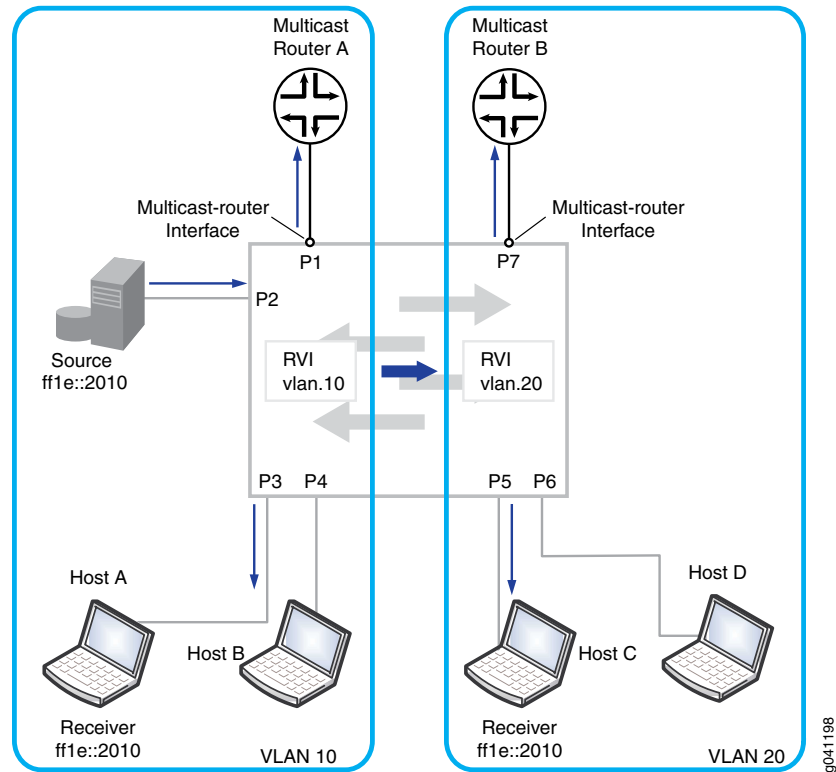


Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs

In the topology shown in [Figure 10 on page 21](#), a multicast source, Multicast Router A, and Hosts A and B are connected to the switch and are in VLAN 10. Multicast Router B and Hosts C and D are also connected to the switch and are in VLAN 20.

In a pure Layer 2 environment, traffic is not forwarded between VLANs. For Host C to receive the multicast traffic from the source on VLAN 10, RVIs must be created on VLAN 10 and VLAN 20 to permit routing of the multicast traffic between the VLANs. In addition, PIM must be enabled on the switch to perform the multicast routing.

Figure 10: Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs



Related Documentation

- [Example: Configuring MLD Snooping on page 30](#)
- [Configuring MLD Snooping on a VLAN \(CLI Procedure\) on page 50](#)
- [Verifying MLD Snooping \(CLI Procedure\) on page 170](#)

CHAPTER 3

Multicast VLAN Registration Overview

- [Understanding Multicast VLAN Registration on page 23](#)

Understanding Multicast VLAN Registration

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

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

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

This topic includes:

- [How MVR Works on page 23](#)

How MVR Works

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

When you configure an MVLAN, you assign a range of multicast group addresses to it. You then configure other VLANs to be MVR receiver VLANs, which receive multicast streams from the MVLAN. The MVR receiver ports comprise all the interfaces that exist

on any of the MVR receiver VLANs. Interfaces that are on the MVLAN itself cannot be MVR receiver ports for that MVLAN.



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

MVR Modes

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

- [MVR Transparent Mode on page 24](#)
- [MVR Proxy Mode on page 24](#)

MVR Transparent Mode

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

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

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

MVR Proxy Mode

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

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

Related Documentation

- [Understanding FIP Snooping, FBF, and MVR Filter Scalability](#)
- [Example: Configuring Multicast VLAN Registration on page 33](#)
- [Configuring Multicast VLAN Registration \(CLI Procedure\) on page 60](#)

PART 2

Configuration

- [Configuration Examples on page 27](#)
- [Configuration Tasks on page 39](#)
- [Configuration Statements on page 63](#)

CHAPTER 4

Configuration Examples

- [Example: Configuring IGMP Snooping on EX Series Switches on page 27](#)
- [Example: Configuring MLD Snooping on page 30](#)
- [Example: Configuring Multicast VLAN Registration on page 33](#)

Example: Configuring IGMP Snooping on EX Series Switches

You can enable IGMP snooping on a VLAN to constrain the flooding of IPv4 multicast traffic on a VLAN. When IGMP snooping is enabled, a switch examines IGMP messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

This example describes how to configure IGMP snooping:

- [Requirements on page 27](#)
- [Overview and Topology on page 28](#)
- [Configuration on page 29](#)
- [Verifying IGMP Snooping Operation on page 29](#)

Requirements

This example uses the following software and hardware components:

- One EX4300 Series switch
- Junos OS Release 13.2 or later for EX Series switches

Before you configure IGMP snooping, be sure you have:

- Configured the **vlan100** VLAN on the switch
- Assigned interfaces **ge-0/0/0**, **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/12** to **vlan100**
- Configure **ge-0/0/12** as a trunk interface.


See <add topic-ref to the Configuring VLANs topic>.

Overview and Topology

In this example, interfaces **ge-0/0/0**, **ge-0/0/1**, and **ge-0/0/2** on the switch are in **vlan100** and are connected to hosts that are potential multicast receivers. Interface **ge-0/0/12**, a trunk interface also in **vlan100**, is connected to a multicast router. The router acts as the IGMP querier and forwards multicast traffic for group **255.100.100.100** to the switch from a multicast source.

The example topology is illustrated in [Figure 11 on page 28](#).

Figure 11: Example IGMP Snooping Topology



In this example topology, the multicast router forwards multicast traffic to the switch from the source when it receives a membership report for group **255.100.100.100** from one of the hosts—for example, Host B. If IGMP snooping is not enabled on **vlan100**, the switch floods the multicast traffic on all interfaces in **vlan100** (except for interface **ge-0/0/12**). If IGMP snooping is enabled on **vlan100**, the switch monitors the IGMP messages between the hosts and router, allowing it to determine that only Host B is interested in receiving the multicast traffic. The switch then forwards the multicast traffic only to interface **ge-0/0/1**.

IGMP snooping is enabled on all VLANs in the default factory configuration. For many implementations, IGMP snooping requires no additional configuration. This example shows how to perform the following optional configurations, which can reduce group join and leave latency:

- Configure immediate leave on the VLAN. When immediate leave is configured, the switch stops forwarding multicast traffic on an interface when it detects that the last member of the multicast group has left the group. If immediate leave is not configured, the switch waits until the group-specific queries time out before it stops forwarding traffic.

Immediate leave is supported by IGMP version 2 (IGMPv2) and IGMPv3. With IGMPv2, we recommend that you configure immediate leave only when there is only one IGMP host on an interface. In IGMPv2, only one host on a interface sends a membership report in response to a group-specific query—any other interested hosts suppress their reports to avoid a flood of reports for the same group. This report-suppression feature means that the switch only knows about one interested host at any given time.

- Configure **ge-0/0/12** as a static multicast-router interface. In this topology, **ge-0/0/12** always leads to the multicast router. By statically configuring **ge-0/0/12** as a multicast-router interface, you avoid any delay imposed by the switch having to learn that **ge-0/0/12** is a multicast-router interface.

Configuration

To configure IGMP snooping on a switch:

CLI Quick Configuration

To quickly configure IGMP snooping, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols igmp-snooping vlan vlan100 immediate-leave
set protocols igmp-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

Step-by-Step Procedure

To configure IGMP snooping on vlan100:

1. Configure the switch to immediately remove a group membership from an interface when it receives a leave report from the last member of the group on the interface:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan100 immediate-leave
```

2. Statically configure interface **ge-0/0/12** as a multicast-router interface:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan100 interface ge-0/0/12
multicast-router-interface
```

Results

Check the results of the configuration:

```
[edit protocols]
user@switch# show igmp-snooping
vlan all;
vlan vlan100 {
    immediate-leave;
    interface ge-0/0/12.0 {
        multicast-router-interface;
    }
}
```

Verifying IGMP Snooping Operation

To verify that IGMP snooping is operating as configured, perform the following task:

- [Displaying IGMP Snooping Information for VLAN vlan100 on page 30](#)

Displaying IGMP Snooping Information for VLAN vlan100

| | |
|------------------------------|---|
| Purpose | Verify that IGMP snooping is enabled on vlan100 and that ge-0/0/12 is recognized as a multicast-router interface. |
| Action | Enter the following command: user@switch> show igmp-snooping vlans vlan vlan100 detail VLAN: vlan100, Tag: 100 Interface: ge-0/0/12.0, tagged, Groups: 0, Router |
| Meaning | By showing information for vlan100 , the command output confirms that IGMP snooping is configured on the VLAN. Interface ge-0/0/12.0 is listed as multicast-router interface, as configured. Because none of the host interfaces are listed, none of the hosts are currently receivers for the multicast group. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IGMP Snooping (CLI Procedure) on page 39• Verifying IGMP Snooping (CLI Procedure) on page 168• IGMP Snooping on EX Series Switches Overview on page 3 |

Example: Configuring MLD Snooping

You can enable MLD snooping on a VLAN to constrain the flooding of IPv6 multicast traffic on a VLAN. When MLD snooping is enabled, a switch examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards IPv6 multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

This example describes how to configure MLD snooping:

- [Requirements on page 30](#)
- [Overview and Topology on page 31](#)
- [Configuration on page 32](#)
- [Verifying MLD Snooping Configuration on page 32](#)

Requirements

This example uses the following software and hardware components:

- One EX Series switch
- Junos OS Release 12.1 or later

Before you configure MLD snooping, be sure you have:

- Configured the **vlan100** VLAN on the switch

- Assigned interfaces **ge-0/0/0**, **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/12** to **vlan100**
- Configured **ge-0/0/12** as a trunk interface.


See *Configuring VLANs for EX Series Switches (CLI Procedure)*.

Overview and Topology

In this example, interfaces **ge-0/0/0**, **ge-0/0/1**, and **ge-0/0/2** on the switch are in **vlan100** and are connected to hosts that are potential multicast receivers. Interface **ge-0/0/12**, a trunk interface also in **vlan100**, is connected to a multicast router. The router acts as the MLD querier and forwards multicast traffic for group **ff1e::2010** to the switch from a multicast source.

The example topology is illustrated in [Figure 12 on page 31](#).

Figure 12: Example MLD Snooping Topology



In this example topology, the multicast router forwards multicast traffic to the switch from the source when it receives a membership report for group **ff1e::2010** from one of the hosts—for example, Host B. If MLD snooping is not enabled on **vlan100**, the switch floods the multicast traffic on all interfaces in **vlan100** (except for interface **ge-0/0/12**). If MLD snooping is enabled on **vlan100**, the switch monitors the MLD messages between the hosts and router, allowing it to determine that only Host B is interested in receiving the multicast traffic. The switch then forwards the multicast traffic only to interface **ge-0/0/1**.

This example shows how to enable MLD snooping on **vlan100**. It also shows how to perform the following optional configurations, which can reduce group join and leave latency:

- Configure immediate leave on the VLAN. When immediate leave is configured, the switch stops forwarding multicast traffic on an interface when it detects that the last

member of the multicast group has left the group. If immediate leave is not configured, the switch waits until the group-specific membership queries time out before it stops forwarding traffic.

- Configure **ge-0/0/12** as a static multicast-router interface. In this topology, **ge-0/0/12** always leads to the multicast router. By statically configuring **ge-0/0/12** as a multicast-router interface, you avoid any delay imposed by the switch having to learn that **ge-0/0/12** is a multicast-router interface.

Configuration

To configure MLD snooping on a switch:

CLI Quick Configuration

To quickly configure MLD snooping, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols mld-snooping vlan vlan100
set protocols mld-snooping vlan vlan100 immediate-leave
set protocols mld-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

Step-by-Step Procedure

To configure MLD snooping:

1. Enable MLD snooping on VLAN **vlan100**:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100
```
2. Configure the switch to immediately remove a group membership from an interface when it receives a leave report from the last member of the group on the interface:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100 immediate-leave
```
3. Statically configure interface **ge-0/0/12** as a multicast-router interface:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

Results

Check the results of the configuration:

```
[edit protocols]
user@switch# show mld-snooping
vlan vlan100 {
    immediate-leave;
    interface ge-0/0/12.0 {
        multicast-router-interface;
    }
}
```

Verifying MLD Snooping Configuration

To verify that MLD snooping is enabled on the VLAN and the MLD snooping forwarding interfaces are correct, perform the following task:

- [Verifying MLD Snooping Interface Membership on VLAN vlan100 on page 33](#)

Verifying MLD Snooping Interface Membership on VLAN **vlan100**

| | |
|------------------------------|---|
| Purpose | Verify that MLD snooping is enabled on vlan100 and that the multicast-router interface is statically configured: |
| Action | <p>Show the group memberships maintained by MLD snooping for vlan100:</p> <pre>user@switch> show mld-snooping membership vlan vlan100 detail VLAN: vlan100 Tag: 100 (Index: 8) Router interfaces: ge-0/0/12.0 static Uptime: 00:15:03 Group: ff1e::2010 ge-0/0/1.0 Timeout: 225 Flags: <V2-hosts> Last reporter: fe80::2020:1:1:3</pre> |
| Meaning | MLD snooping is running on vlan100 , and interface ge-0/0/12.0 is a statically configured multicast-router interface. Because the multicast group ff1e::2010 is listed, at least one host in the VLAN is a current member of the multicast group and that host is on interface ge-0/0/1.0 . |
| Related Documentation | <ul style="list-style-type: none"> • Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 • Verifying MLD Snooping (CLI Procedure) on page 170 • Understanding MLD Snooping on page 13 |

Example: Configuring Multicast VLAN Registration

Multicast VLAN registration (MVR) enables hosts that are not part of a multicast VLAN (MVLAN) to receive multicast streams from the MVLAN, which enable the MVLAN to be shared across the Layer 2 network and eliminate the need to send duplicate multicast streams to each requesting VLAN in the network. Hosts remain in their own VLANs for bandwidth and security reasons.

This example describes how to configure MVR on EX Series switches and the QFX Series.

- [Requirements on page 33](#)
- [Overview and Topology on page 34](#)
- [Configuration on page 36](#)

Requirements

This example uses the following hardware and software components:

- One EX Series or QFX Series switch
- Junos OS Release 9.6 or later for EX Series switches or Junos OS Release 12.3 or later for the QFX Series

Before you configure MVR, be sure you have:

- Configured two or more VLANs on the switch. See the task for your platform:

- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- *Example: Setting Up Bridging with Multiple VLANs for the QFX Series and EX4600 switch*
- Connected the switch to a network that can transmit IPTV multicast streams from a video server.
- Connected a host that is capable of receiving IPTV multicast streams to an interface in one of the VLANs.

Overview and Topology

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

MVR introduces the concept of a *multicast source VLAN* (MVLAN), which is created by MVR and becomes the only VLAN over which multicast traffic flows throughout the Layer 2 network. Multicast traffic can then be selectively forwarded from interfaces on the MVLAN (source ports) to hosts that are connected to interfaces (multicast receiver ports) that are not part of the multicast source VLAN. When you configure an MVLAN, you assign a range of multicast group addresses to it. You then configure other VLANs to be MVR receiver VLANs, which receive multicast streams from the MVLAN. The MVR receiver ports comprise all the interfaces that exist on any of the MVR receiver VLANs.

You can configure MVR to operate in one of two modes: transparent mode (the default mode) or proxy mode. Both modes enable MVR to forward only one copy of a multicast stream to the Layer 2 network.

In transparent mode, the switch receives one copy of each IPTV multicast stream and then replicates the stream only to those hosts that want to receive it, while forwarding all other types of multicast traffic without modification. [Figure 13 on page 35](#) shows how MVR operates in transparent mode.

In proxy mode, the switch acts as a proxy for the IGMP multicast router in the MVLAN for MVR group memberships established in the MVR receiver VLANs and generates and sends IGMP packets into the MVLAN as needed. [Figure 14 on page 36](#) shows how MVR operates in proxy mode.

This example shows how to configure MVR in both transparent mode and proxy mode on an EX Series switch or the QFX Series. The topology includes a video server that is connected to a multicast router, which in turn forwards the IPTV multicast traffic in the MVLAN to the Layer 2 network.

[Figure 13 on page 35](#) shows the MVR topology in transparent mode. Interfaces P1 and P2 on Switch C belong to service VLAN s0 and MVLAN mv0. Interface P4 of Switch C also belongs to service VLAN s0. In the upstream direction of the network, only non-IPTV traffic is being carried in individual customer VLANs of service VLAN s0. VLAN c0 is an example of this type of customer VLAN. IPTV traffic is being carried on MVLAN mv0. If any host on any customer VLAN connected to port P4 requests an MVR stream, Switch C

takes the stream from VLAN mv0 and replicates that stream onto port P4 with tag mv0. IPTV traffic, along with other network traffic, flows from port P4 out to the Digital Subscriber Line Access Multiplexer (DSLAM) D1.

Figure 13: MVR Topology in Transparent Mode





Figure 14 on page 36 shows the MVR topology in proxy mode. Interfaces P1 and P2 on Switch C belong to MVLAN mv0 and customer VLAN c0. Interface P4 on Switch C is an access port of customer VLAN c0. In the upstream direction of the network, only non-IPTV traffic is being carried on customer VLAN c0. Any IPTV traffic requested by hosts on VLAN c0 is replicated untagged to port P4 based on streams received in MVLAN mv0. IPTV traffic flows from port P4 out to an IPTV-enabled device in Host H1. Other traffic, such as data and voice traffic, also flows from port P4 to other network devices in Host H1.

Figure 14: MVR Topology in Proxy Mode



For information on VLAN tagging, see the topic for your platform:

- *Understanding Bridging and VLANs on EX Series Switches*
- *Understanding Bridging and VLANs on the QFX Series and EX4600 switch*

Configuration

CLI Quick Configuration

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

```
set vlan mv0 data-forwarding source groups 225.10.0.0/16
set vlan v2 data-forwarding receiver source-vlans mv0
set vlan v2 data-forwarding receiver install
set vlan mv0 proxy source-address 10.1.1.1
```

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

To configure MVR:

1. Configure VLAN mv0 to be an MVLAN:

```
[edit protocols igmp-snooping]
user@switch# set vlan mv0 data-forwarding source groups 225.10.0.0/16
```
2. Configure VLAN v2 to be a multicast receiver VLAN with mv0 as its source:

```
[edit protocols igmp-snooping]
user@switch# set vlan v2 data-forwarding receiver source-vlans mv0
```
3. (Optional) Install forwarding entries in the multicast receiver VLAN v2:

```
[edit protocols igmp-snooping]
user@switch# set vlan v2 data-forwarding receiver install
```
4. (Optional) Configure MVR in proxy mode:

```
[edit protocols igmp-snooping]
user@switch# set vlan mv0 proxy source-address 10.1.1.1
```

Results From configuration mode, confirm your configuration by entering the **show** command at the **[edit protocols igmp-snooping]** hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit protocols igmp-snooping]
user@switch# show
vlan mv0 {
  proxy {
    source-address 10.1.1.1;
  }
  data-forwarding {
    source {
      groups 225.10.0.0/16;
    }
  }
}
vlan v2 {
  data-forwarding {
    receiver {
      source-vlans mv0;
      install;
    }
  }
}
```

Related Documentation

- [Configuring Multicast VLAN Registration \(CLI Procedure\) on page 60](#)
- [Understanding Multicast VLAN Registration on page 23](#)

CHAPTER 5

Configuration Tasks

- [Configuring IGMP Snooping \(CLI Procedure\) on page 39](#)
- [Configuring IGMP Snooping \(J-Web Procedure\) on page 46](#)
- [Configuring IGMP Snooping Tracing Operations \(CLI Procedure\) on page 48](#)
- [Configuring MLD Snooping on a VLAN \(CLI Procedure\) on page 50](#)
- [Configuring MLD Snooping Tracing Operations \(CLI Procedure\) on page 58](#)
- [Configuring Multicast VLAN Registration \(CLI Procedure\) on page 60](#)

Configuring IGMP Snooping (CLI Procedure)

IGMP snooping constrains the flooding of IPv4 multicast traffic on a VLAN. When IGMP snooping is enabled, a switch examines IGMP messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

The factory default configuration enables IGMP snooping on all VLANs. For many networks, IGMP snooping requires no further configuration.

You can perform the following optional configurations per VLAN:

- Selectively enable IGMP snooping on specific VLANs.



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

- Specify the IGMP version for the general query that the switch sends on an interface when the interface comes up.
- Enable immediate leave on a VLAN or all VLANs. Immediate leave reduces the length of time it takes the switch to stop forwarding multicast traffic when the last member host on the interface leaves the group.
- Configure an interface as a static multicast-router interface for a VLAN or for all VLANs so that the switch does not need to dynamically learn that the interface is a multicast-router interface.

- Configure an interface as a static member of a multicast group so that the switch does not need to dynamically learn the interface's membership.
- Change the value for certain timers and counters to match the values configured on the multicast router serving as the IGMP querier.
- Configure multicast VLAN registration (MVR). MVR allows hosts that are not part of a multicast source VLAN (MVLAN) to still receive multicast streams from the MVLAN, allowing an MVLAN to be shared across a Layer 2 network. See [“Configuring Multicast VLAN Registration \(CLI Procedure\)” on page 60](#) for more information.



TIP: When you configure IGMP snooping using the `vlan all` statement, any VLAN that is not individually configured for IGMP snooping inherits the `vlan all` configuration. Any VLAN that is individually configured for IGMP snooping, on the other hand, inherits none of its configuration from `vlan all`. Any parameters that are not explicitly defined for the individual VLAN assume their default values, not the values specified in the `vlan all` configuration. For example, in the following configuration:

```
protocols {
  igmp-snooping {
    vlan all {
      robust-count 8;
    }
    vlan employee {
      interface ge-0/0/8.0 {
        static {
          group 225.0.0.1;
        }
      }
    }
  }
}
```

all VLANs, except `employee`, have a robust count of 8. Because `employee` has been individually configured, its robust count value is not determined by the value set under `vlan all`. Instead, its robust count is the default value of 2.

This topic covers:

- [Enabling or Disabling IGMP Snooping on VLANs on page 41](#)
- [Configuring the IGMP Version on page 41](#)
- [Enabling Immediate Leave on page 42](#)
- [Configuring an Interface as a Multicast-Router Interface on page 43](#)
- [Configuring Static Group Membership on an Interface on page 44](#)
- [Changing the Timer and Counter Values on page 45](#)

Enabling or Disabling IGMP Snooping on VLANs

The factory default configuration on EX Series switches enables IGMP snooping on all VLANs by including the following configuration:

```
protocols {
  igmp-snooping {
    vlan all;
  }
}
```

This topic describes how you can selectively enable or disable IGMP snooping on VLANs. It assumes you are starting from the factory default configuration.

- To disable IGMP snooping on a VLAN:

```
[edit protocols igmp-snooping]
user@switch# set vlan vlan-name disable
```

- To enable IGMP snooping on only a few VLANs:

1. Disable IGMP snooping on all VLANs:

```
[edit protocols igmp-snooping]
user@switch# set vlan all disable
```

2. Enable IGMP snooping on each individual VLAN:

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

For example, to enable IGMP snooping only on VLANs **sales** and **support**:

```
[edit protocols igmp-snooping]
user@switch# set vlan all disable
```

```
[edit protocols igmp-snooping]
user@switch# set vlan sales
```

```
[edit protocols igmp-snooping]
user@switch# set vlan support
```

You can also deactivate the IGMP snooping protocol on the switch without changing the IGMP snooping VLAN configurations:

```
[edit]
user@switch# deactivate protocols igmp-snooping
```

Configuring the IGMP Version

You can configure the version of IGMP queries sent by a switch when IGMP snooping is enabled. By default, the switch uses IGMP version 2 (IGMPv2). If you are using Protocol-Independent Multicast source-specific Multicast (PIM-SSM), we recommend that you configure the switch to use IGMPv3.

Typically, a switch passively monitors IGMP messages sent between multicast routers and hosts and does not send IGMP queries. The exception is when a switch detects that an interface has come up. When an interface comes up, the switch sends an immediate general membership query to all hosts on the interface. By doing so, the switch enables

the multicast routers to learn group memberships more quickly than they would if they had to wait until the IGMP querier sent its next general query.

The IGMP version of the general query determines the IGMP version of the host membership reports as follows:

- IGMP version 1 (IGMPv1) general query—IGMPv1, IGMPv2, and IGMPv3 hosts respond with an IGMPv1 membership report.
- IGMPv2 general query—IGMPv2 and IGMPv3 hosts respond with an IGMPv2 membership report, while IGMPv1 hosts respond with a IGMPv1 membership report.
- IGMPv3 general query—IGMPv3 hosts respond with an IGMPv3 membership report, while IGMPv1 and IGMPv2 hosts are unable to respond to the query.

By default, the switch sends IGMPv2 queries. If your VLAN contains IGMPv3 multicast routers and hosts and the routers are running PIM-SSM, we recommend that you configure IGMP snooping for IGMPv3. Doing so enables the routers to quickly learn which multicast sources the hosts on the interface want to receive traffic from.



NOTE: Configuring the IGMP version does not limit the version of IGMP messages that the switch can snoop. A switch can snoop both IGMPv1, IGMPv2,, and IGMPv3 messages regardless of the IGMP version configured.

To configure the IGMP version on a switch:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name version number
```

For example, to set the IGMP version to version 3 for VLAN **marketing**:

```
[edit protocols]
user@switch# set igmp-snooping vlan marketing version 3
```

Enabling Immediate Leave

By default, when a switch with IGMP snooping enabled receives an IGMP leave report on a member interface, it waits for hosts on the interface to respond to IGMP group-specific queries to determine whether there still are hosts on the interface interested in receiving the group multicast traffic. If the switch does not see any membership reports for the group within a set interval of time, it removes the interface's group membership from the multicast forwarding table and stops forwarding multicast traffic for the group to the interface.

You can decrease the leave latency created by this default behavior by enabling immediate leave on a VLAN.

When you enable immediate leave on a VLAN, host tracking is also enabled, allowing the switch to keep track of the hosts on a interface that have joined a multicast group. When the switch receives a leave report from the last member of the group, it immediately stops forwarding traffic to the interface and does not wait for the interface group membership to time out.

Immediate leave is supported for both IGMP version 2 (IGMPv2) and IGMPv3. However, with IGMPv2, we recommend that you configure immediate leave only when there is only one IGMP host on an interface. In IGMPv2, only one host on a interface sends a membership report in response to a group-specific query—any other interested hosts suppress their reports to avoid a flood of reports for the same group. This report-suppression feature means that the switch only knows about one interested host at any given time.

To enable immediate leave on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name immediate-leave
```

To enable immediate leave on all VLANs:

```
[edit protocols]
user@switch# set igmp-snooping vlan all immediate-leave
```

Configuring an Interface as a Multicast-Router Interface

When IGMP snooping is enabled on a switch, the switch determines which interfaces face a multicast router by monitoring interfaces for IGMP queries or Protocol Independent Multicast (PIM) updates. If the switch receives these messages on an interface, it adds the interface to its multicast forwarding table as a multicast-router interface.

In addition to dynamically learned interfaces, the multicast forwarding table can include interfaces that you explicitly configure to be multicast router interfaces. Unlike the table entries for dynamically learned interfaces, table entries for statically configured interfaces are not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure a static multicast-router interface include:

- You have an unusual network configuration that prevents IGMP snooping from reliably learning about a multicast-router interface through monitoring IGMP queries or PIM updates.
- Your implementation does not require an IGMP querier.
- You have a stable topology and want to avoid the delay the dynamic learning process entails.



NOTE: If the interface you are configuring as a multicast-router interface is a trunk port, the interface becomes a multicast-router interface for all VLANs configured on the trunk port even if you have not explicitly configured it for all the VLANs. In addition, all unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded to the multicast-router interface, even if the interface is configured as a multicast-router interface only for IGMP snooping.

To configure an interface as a static multicast-router interface:

```
[edit protocols]
```

```
user@switch# set igmp-snooping vlan vlan-name interface interface-name
multicast-router-interface
```

For example, to configure **ge-0/0/5.0** as a multicast-router interface for all VLANs on the switch:

```
[edit protocols]
user@switch# set igmp-snooping vlan all interface ge-0/0/5.0
multicast-router-interface
```

Configuring Static Group Membership on an Interface

To determine how to forward multicast packets, a switch with IGMP snooping enabled maintains a multicast forwarding table containing a list of host interfaces that have interested listeners for a specific multicast group. The switch learns which host interfaces to add or delete from this table by examining IGMP membership reports as they arrive on interfaces on which IGMP snooping is enabled.

In addition to such dynamically learned interfaces, the multicast forwarding table can include interfaces that you statically configure to be members of multicast groups. When you configure a static group interface, the switch adds the interface to the forwarding table as a host interface for the group. Unlike an entry for a dynamically learned interface, a static interface entry is not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure static group membership on an interface include:

- You want to simulate an attached multicast receiver for testing purposes.
- The interface has receivers that cannot send IGMP membership reports.
- You want the multicast traffic for a specific group to be immediately available to a receiver without any delay imposed by the dynamic join process.

You cannot configure multicast source addresses for a static group interface.



NOTE: The switch does not simulate IGMP membership reports on behalf of a statically configured interface. Thus a multicast router might be unaware that the switch has an interface that is a member of the multicast group. You can configure a static group interface on the router to ensure that the switch receives the group multicast traffic.

To configure a host interface as a static member of a multicast group:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name interface interface-name static group
ip-address
```

For example, to configure interface **ge-0/0/11.0** in VLAN **ip-camera-vlan** as a static member of multicast group **225.0.0.1**:

```
[edit protocols]
```

```
user@switch# set igmp-snooping vlan ip-camera-vlan interface ge-0/0/11.0 static
group 225.0.0.1
```

Changing the Timer and Counter Values

IGMP uses various timers and counters to determine how often an IGMP querier sends out membership queries and when group memberships time out. On Juniper Networks EX Series Ethernet Switches, the IGMP and IGMP snooping timers and counters default values are set to the values recommended in RFC 2236, *Internet Group Management Protocol, Version 2*. These values work well for most multicast implementations.

There might be cases, however, where you might want to adjust the timer and counter values—for example, to reduce burstiness, to reduce leave latency, or to adjust for expected packet loss on a subnet. If you change a timer or counter value for the IGMP querier on a VLAN, we recommend that you change the value for all multicast routers and switches on the VLAN so that all devices time out group memberships at approximately the same time.

The following timers and counters are configurable on a switch:

- **query-interval**—The length of time the IGMP querier waits between sending general queries (the default is 125 seconds). You can change this interval to tune the number of IGMP messages on the subnet; larger values cause general queries to be sent less often.

You cannot configure this value directly for IGMP snooping. IGMP snooping inherits the value from the IGMP value configured on the switch, which is applied to all VLANs on the switch.

To configure the IGMP **query-interval**:

```
[edit protocols]
user@switch# set igmp query-interval seconds
```

- **query-response-interval**—The maximum length of time the host can wait until it responds (the default is 10 seconds). You can change this interval to adjust the burst peaks of IGMP messages on the subnet. Set a larger interval to make the traffic less bursty.

You cannot configure this value directly for IGMP snooping. IGMP snooping inherits the value from the IGMP value configured on the switch, which is applied to all VLANs on the switch.

To configure the IGMP **query-response-interval**:

```
[edit protocols]
user@switch# set igmp query-response-interval seconds
```

- **query-last-member-interval**—The length of time the IGMP querier waits between sending group-specific membership queries (the default is 1 second). The IGMP querier sends a group-specific query after receiving a leave report from a host. You can decrease this interval to reduce the amount of time it takes for multicast traffic to stop forwarding after the last member leaves a group.

You cannot configure this value directly for IGMP snooping. IGMP snooping inherits the value from the IGMP value configured on the switch, which is applied to all VLANs on the switch.

To configure the IGMP **query-last-member-interval**:

```
[edit protocols]
user@switch# set igmp query-last-member-interval seconds
```

- **robust-count**—The number of times the querier resends a general membership query or a group-specific membership query (the default is 2 times). You can increase this count to tune for higher expected packet loss.

For IGMP snooping, you can configure **robust-count** for a specific VLAN. If a VLAN does not have **robust-count** configured, the **robust-count** value is inherited from the value configured for IGMP.

To configure **robust-count** for IGMP snooping on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name robust-count number
```

The values configured for **query-interval**, **query-response-interval**, and **robust-count** determine the multicast listener interval—the length of time the switch waits for a group membership report after a general query before removing a multicast group from its multicast forwarding table. The switch calculates the multicast listener interval by multiplying **query-interval** by **robust-count** and then adding **query-response-interval**:

$(\text{query-interval} \times \text{robust-count}) + \text{query-response-interval} = \text{multicast listener interval}$

For example, the multicast listener interval is 260 seconds when the default settings for **query-interval**, **query-response-interval**, and **robust-count** are used:

$(125 \times 2) + 10 = 260$

You can display the time remaining in the multicast listener interval before a group times out by using the **show igmp-snooping membership** command.

Related Documentation

- [Example: Configuring IGMP Snooping on EX Series Switches on page 27](#)
- [Example: Configuring Multicast VLAN Registration on page 33](#)
- [Verifying IGMP Snooping \(CLI Procedure\) on page 168](#)
- [IGMP Snooping on EX Series Switches Overview on page 3](#)
- [Understanding Multicast VLAN Registration on page 23](#)

Configuring IGMP Snooping (J-Web Procedure)



NOTE: This topic applies only to the J-Web Application package.

IGMP snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, the EX Series switch monitors the IGMP transmissions between a host (a network

device) and a multicast router, keeping track of the multicast groups and associated member interfaces. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces.

You can configure IGMP snooping on one or more VLANs to allow the switch to examine IGMP packets and make forwarding decisions based on packet content. By default, IGMP snooping is enabled on EX Series switches.

To enable IGMP snooping and configure individual options by using the J-Web interface:

1. Select **Configure > Switching > IGMP Snooping**.



NOTE: After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—Creates an IGMP snooping configuration for the VLAN.
- **Edit**—Modifies an IGMP snooping configuration for the VLAN.
- **Delete**—Deletes a selected VLAN from the IGMP snooping configuration.

When you are adding or editing an IGMP snooping configuration, enter information as described in [Table 3 on page 47](#).

3. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.

To disable IGMP snooping on a VLAN, select the VLAN from the list and click **Disable**.



NOTE: The **Disable** option is not available for EX4300 switches.

Table 3: IGMP Snooping Configuration Fields

| Field | Function | Your Action |
|-----------------|---|--|
| VLAN Name | Specifies the VLAN on which to enable IGMP snooping. | Select a VLAN from the list to add it to the snooping configuration. |
| Immediate Leave | Immediately removes a multicast group membership from an interface when it receives a leave message from that interface without waiting for any other IGMP messages to be exchanged (IGMP version 2 and IGMP version 3 only). | To enable the option, select the check box. To disable the option, clear the check box. |

Table 3: IGMP Snooping Configuration Fields (*continued*)

| Field | Function | Your Action |
|-----------------|--|---|
| Robust Count | Specifies the number of timeout intervals the switch waits before timing out a multicast group. | Type a value. |
| Interfaces List | Statically configures an interface as a switching interface toward a multicast router or as a member of a multicast group. | <p>Click one of the following options:</p> <ul style="list-style-type: none"> • Add—Adds an interface to the IGMP snooping configuration. <ol style="list-style-type: none"> 1. Select an interface from the list. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list. 2. Select Multicast Router Interface. 3. Type the maximum number of groups an interface can join. 4. In Static, choose one: <ul style="list-style-type: none"> • Click Add, type a group IP address, and click OK. • Select a group and click Remove to remove the group membership. • Edit—Edits the interface settings for the IGMP snooping configuration. • Remove—Deletes an interface configured for IGMP snooping. |

Related Documentation

- [Example: Configuring IGMP Snooping on EX Series Switches on page 27](#)
- [Configuring IGMP Snooping \(CLI Procedure\) on page 39](#)
- [IGMP Snooping on EX Series Switches Overview on page 3](#)

Configuring IGMP Snooping Tracing Operations (CLI Procedure)

By enabling tracing operations for IGMP snooping, you can record detailed messages about the operation of the protocol, such as the various types of protocol packets sent and received. [Table 4 on page 48](#) describes the tracing operations you can enable and the flags used to specify them in the tracing configuration.

Table 4: Supported Tracing Operations for IGMP Snooping

| Tracing Operation | Flag |
|--|----------------|
| Trace all (equivalent of including all flags). | all |
| Trace general IGMP snooping protocol events. | general |

Table 4: Supported Tracing Operations for IGMP Snooping (*continued*)

| Tracing Operation | Flag |
|--|----------------|
| Trace communication over routing socket events. | krt |
| Trace leave reports (IGMPv2 and IGMPv3 only). | leave |
| Trace nexthop-related events. | nexthop |
| Trace normal IGMP snooping protocol events. If you do not specify this flag, only unusual or abnormal operations are traced. | normal |
| Trace all IGMP packets. | packets |
| Trace policy processing. | policy |
| Trace IGMP membership query messages. | query |
| Trace membership reports | report |
| Trace routing information. | route |
| Trace state transitions. | state |
| Trace routing protocol task processing. | task |
| Trace timer processing. | timer |
| Trace VLAN-related events. | vlan |

This topic covers:

- [Configuring Tracing Operations on page 49](#)
- [Viewing, Stopping, and Restarting Tracing Operations on page 50](#)

Configuring Tracing Operations

To configure tracing operations for IGMP snooping:

1. Configure the filename for the trace file:

```
[edit protocols igmp-snooping ]
user@switch# set traceoptions file filename
```

For example:

```
[edit protocols igmp-snooping ]
user@switch# set traceoptions file mld-snoop-trace
```

2. (Optional) Configure the maximum number of trace files and size of the trace files:

```
[edit protocols igmp-snooping ]
user@switch # set file files number size size
```

For example:

```
[edit protocols igmp-snooping ]
user@switch # set traceoptions file files 5 size 1m
```

causes the contents of the trace file to be emptied and archived in a .gz file when the file reaches 1 MB. Four archive files are maintained, the contents of which are rotated whenever the current active trace file is archived.

If you omit this step, the maximum number of trace files defaults to 10, with the maximum file size defaulting to 128 K.

3. Specify one of the tracing flags shown in [Table 4 on page 48](#):

```
[edit protocols igmp-snooping ]
user@switch # set traceoptions flag flagname
```

For example, to perform trace operations on VLAN-related events and IGMP query messages:

```
[edit protocols igmp-snooping ]
user@switch# set traceoptions flag vlan

[edit protocols igmp-snooping ]
user@switch# set traceoptions flag query
```

Viewing, Stopping, and Restarting Tracing Operations

When you commit the configuration, tracing operations begin. You can view the trace file in the `/var/log` directory. For example:

```
user@switch> file show /var/log/igmp-snoop-trace
```

You can stop and restart tracing operations by deactivating and reactivating the configuration:

```
[edit]
user@switch# deactivate protocols igmp-snooping traceoptions

[edit]
user@switch# activate protocols igmp-snooping traceoptions
```

- Related Documentation**
- [Configuring IGMP Snooping \(CLI Procedure\) on page 39](#)
 - *Tracing and Logging Junos OS Operations*

Configuring MLD Snooping on a VLAN (CLI Procedure)

You can enable MLD snooping on a VLAN to constrain the flooding of IPv6 multicast traffic on a VLAN. When MLD snooping is enabled, a switch examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards IPv6 multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

MLD snooping is not enabled on the switch by default. To enable MLD snooping on all VLANs:

```
[edit]
user@switch# set protocols mld-snooping vlan all
```

For many networks, MLD snooping requires no further configuration.

You can perform the following optional configurations per VLAN:

- Selectively enable MLD snooping on specific VLANs.



NOTE: You cannot configure MLD snooping on a secondary VLAN.

- Specify the MLD version for the general query that the switch sends on an interface when the interface comes up.
- Enable immediate leave on a VLAN or all VLANs. Immediate leave reduces the length of time it takes the switch to stop forwarding multicast traffic when the last member host on the interface leaves the group.
- Configure an interface as a static multicast-router interface for a VLAN or for all VLANs so that the switch does not need to dynamically learn that the interface is a multicast-router interface.
- Configure an interface as a static member of a multicast group so that the switch does not need to dynamically learn the interface's membership.
- Change the value for certain timers and counters to match the values configured on the multicast router serving as the MLD querier.



TIP: When you configure MLD snooping using the `vlan all` statement, any VLAN that is not individually configured for MLD snooping inherits the `vlan all` configuration. Any VLAN that is individually configured for MLD snooping, on the other hand, inherits none of its configuration from `vlan all`. Any parameters that are not explicitly defined for the individual VLAN assume their default values, not the values specified in the `vlan all` configuration. For example, in the following configuration:

```
protocols {
  mld-snooping {
    vlan all {
      robust-count 8;
    }
    vlan employee {
      interface ge-0/0/8.0 {
        static {
          group ff1e::1;
        }
      }
    }
  }
}
```

all VLANs, except `employee`, have a robust count of 8. Because `employee` has been individually configured, its robust count value is not determined by the value set under `vlan all`. Instead, its robust count is the default value of 2.

.....

This topic covers:

- [Enabling or Disabling MLD Snooping on VLANs on page 52](#)
- [Configuring the MLD Version on page 53](#)
- [Enabling Immediate Leave on page 54](#)
- [Configuring an Interface as a Multicast-Router Interface on page 54](#)
- [Configuring Static Group Membership on an Interface on page 55](#)
- [Changing the Timer and Counter Values on page 56](#)

Enabling or Disabling MLD Snooping on VLANs

MLD snooping is not enabled on any VLAN by default. You must explicitly configure a VLAN or all VLANs for MLD snooping.

This topic describes how you can enable or disable MLD snooping on specific VLANs or on all VLANs on the switch.

- To enable MLD snooping on all VLANs:

```
[edit protocols mld-snooping]
user@switch# set vlan all
```

- To enable MLD snooping on a specific VLAN:

```
[edit protocols mld-snooping]
user@switch# set vlan vlan-name
```



NOTE: You cannot configure MLD snooping on a secondary VLAN.

.....

For example, to enable MLD snooping on VLAN `education`:

```
[edit protocols mld-snooping]
user@switch# set vlan education
```

- To enable MLD snooping on all VLANs except a few VLANs:

1. Enable MLD snooping on all VLANs:

```
[edit protocols mld-snooping]
user@switch# set vlan all
```

2. Disable MLD snooping on individual VLANs:

```
[edit protocols mld-snooping]
user@switch# set vlan vlan-name disable
```

For example, to enable MLD snooping on all VLANs except `vlan100` and `vlan200`:

```
[edit protocols mld-snooping]
```

```

user@switch# set vlan all
[edit protocols mld-snooping]
user@switch# set vlan vlan100 disable
[edit protocols mld-snooping]
user@switch# set vlan vlan200 disable

```

You can also deactivate the MLD snooping protocol on the switch without changing the MLD snooping VLAN configurations:

```

[edit]
user@switch# deactivate protocols mld-snooping

```

Configuring the MLD Version

You can configure the version of MLD queries sent by a switch when MLD snooping is enabled. By default, the switch uses MLD version 1 (MLDv1). If you are using Protocol-Independent Multicast source-specific multicast (PIM-SSM), we recommend that you configure the switch to use MLDv2.

Typically, a switch passively monitors MLD messages sent between multicast routers and hosts and does not send MLD queries. The exception is when a switch detects that an interface has come up. When an interface comes up, the switch sends an immediate general membership query to all hosts on the interface. By doing so, the switch enables the multicast routers to learn group memberships more quickly than they would if they had to wait until the MLD querier sent its next general query.

The MLD version of the general query determines the MLD version of the host membership reports as follows:

- MLD version 1 (MLDv1) general query—Both MLDv1 and MLDv2 hosts respond with an MLDv1 membership report.
- MLDv2 general query—MLDv2 hosts respond with an MLDv2 membership report, while MLDv1 hosts are unable to respond to the query.

By default, the switch sends MLDv1 queries. This ensures compatibility with hosts and multicast routers that support MLDv1 only and cannot process MLDv2 reports. However, if your VLAN contains MLDv2 multicast routers and hosts and the routers are running PIM-SSM, we recommend that you configure MLD snooping for MLDv2. Doing so enables the routers to quickly learn which multicast sources the hosts on the interface want to receive traffic from.



NOTE: Configuring the MLD version does not limit the version of MLD messages that the switch can snoop. A switch can snoop both MLDv1 and MLDv2 messages regardless of the MLD version configured.

To configure the MLD version on a switch:

```

[edit protocols]
user@switch# set mld-snooping vlan vlan-name version number

```

For example, to set the MLD version to version 2 for VLAN **marketing**:

```
[edit protocols]
user@switch# set mld-snooping vlan marketing version 2
```

Enabling Immediate Leave

By default, when a switch with MLD snooping enabled receives an MLD leave report on a member interface, it waits for hosts on the interface to respond to MLD group-specific queries to determine whether there still are hosts on the interface interested in receiving the group multicast traffic. If the switch does not see any membership reports for the group within a set interval of time, it removes the interface's group membership from the multicast forwarding table and stops forwarding multicast traffic for the group to the interface.

You can decrease the leave latency created by this default behavior by enabling immediate leave on a VLAN.

When you enable immediate leave on a VLAN, host tracking is also enabled, allowing the switch to keep track of the hosts on a interface that have joined a multicast group. When the switch receives a leave report from the last member of the group, it immediately stops forwarding traffic to the interface and does not wait for the interface group membership to time out.

Immediate leave is supported for both MLD version 1 (MLDv1) and MLDv2. However, with MLDv1, we recommend that you configure immediate leave only when there is only one MLD host on an interface. In MLDv1, only one host on a interface sends a membership report in response to a group-specific query—any other interested hosts suppress their reports. This report-suppression feature means that the switch only knows about one interested host at any given time.

To enable immediate leave on a VLAN:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan-name immediate-leave
```

To enable immediate leave on all VLANs:

```
[edit protocols]
user@switch# set mld-snooping vlan all immediate-leave
```

Configuring an Interface as a Multicast-Router Interface

When MLD snooping is enabled on a switch, the switch determines which interfaces face a multicast router by monitoring interfaces for MLD queries or Protocol Independent Multicast (PIM) updates. If the switch receives these messages on an interface, it adds the interface to its multicast forwarding table as a multicast-router interface.

In addition to dynamically learned interfaces, the multicast forwarding table can include interfaces that you explicitly configure to be multicast router interfaces. Unlike the table entries for dynamically learned interfaces, table entries for statically configured interfaces are not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure a static multicast-router interface include:

- You have an unusual network configuration that prevents MLD snooping from reliably learning about a multicast-router interface through monitoring MLD queries or PIM updates.
- Your implementation does not require an MLD querier.
- You have a stable topology and want to avoid the delay the dynamic learning process entails.



NOTE: If the interface you are configuring as a multicast-router interface is a trunk port, the interface becomes a multicast-router interface for all VLANs configured on the trunk port even if you have not explicitly configured it for all the VLANs. In addition, all unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded to the multicast-router interface, even if the interface is configured as a multicast-router interface only for MLD snooping.

To configure an interface as a static multicast-router interface:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan-name interface interface-name
multicast-router-interface
```

For example, to configure **ge-0/0/5.0** as a multicast-router interface for all VLANs on the switch:

```
[edit protocols]
user@switch# set mld-snooping vlan all interface ge-0/0/5.0 multicast-router-interface
```

Configuring Static Group Membership on an Interface

To determine how to forward multicast packets, a switch with MLD snooping enabled maintains a multicast forwarding table containing a list of host interfaces that have interested listeners for a specific multicast group. The switch learns which host interfaces to add or delete from this table by examining MLD membership reports as they arrive on interfaces on which MLD snooping is enabled.

In addition to such dynamically learned interfaces, the multicast forwarding table can include interfaces that you statically configure to be members of multicast groups. When you configure a static group interface, the switch adds the interface to the forwarding table as a host interface for the group. Unlike an entry for a dynamically learned interface, a static interface entry is not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure static group membership on an interface include:

- You want to simulate an attached multicast receiver for testing purposes.
- The interface has receivers that cannot send MLD membership reports.

- You want the multicast traffic for a specific group to be immediately available to a receiver without any delay imposed by the dynamic join process.

You cannot configure multicast source addresses for a static group interface. The MLD version of a static group interface is always MLD version 1.



NOTE: The switch does not simulate MLD membership reports on behalf of a statically configured interface. Thus a multicast router might be unaware that the switch has an interface that is a member of the multicast group. You can configure a static group interface on the router to ensure that the switch receives the group multicast traffic.

To configure a host interface as a static member of a multicast group:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan-name interface interface-name static group
ip-address
```

For example, to configure interface **ge-0/0/11.0** in VLAN **ip-camera-vlan** as a static member of multicast group **ff1e::1**:

```
[edit protocols]
user@switch# set mld-snooping vlan ip-camera-vlan interface ge-0/0/11.0 static group
ff1e::1
```

Changing the Timer and Counter Values

MLD uses various timers and counters to determine how often an MLD querier sends out membership queries and when group memberships time out. On Juniper Networks EX Series switches, the MLD and MLD snooping timers and counters default values are set to the values recommended in RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*. These values work well for most multicast implementations.

There might be cases, however, where you might want to adjust the timer and counter values—for example, to reduce burstiness, to reduce leave latency, or to adjust for expected packet loss on a subnet. If you change a timer or counter value for the MLD querier on a VLAN, we recommend that you change the value for all multicast routers and switches on the VLAN so that all devices time out group memberships at approximately the same time.

The following timers and counters are configurable on a switch:

- **query-interval**—The length of time the MLD querier waits between sending general queries (the default is 125 seconds). You can change this interval to tune the number of MLD messages on the subnet; larger values cause general queries to be sent less often.

You cannot configure this value directly for MLD snooping. MLD snooping inherits the value from the MLD value configured on the switch, which is applied to all VLANs on the switch.

To configure the MLD **query-interval**:

```
[edit protocols]
user@switch# set mld query-interval seconds
```

- **query-response-interval**—The maximum length of time the host can wait until it responds (the default is 10 seconds). You can change this interval to adjust the burst peaks of MLD messages on the subnet. Set a larger interval to make the traffic less bursty.

You cannot configure this value directly for MLD snooping. MLD snooping inherits the value from the MLD value configured on the switch, which is applied to all VLANs on the switch.

To configure the MLD **query-response-interval**:

```
[edit protocols]
user@switch# set mld query-response-interval seconds
```

- **query-last-member-interval**—The length of time the MLD querier waits between sending group-specific membership queries (the default is 1 second). The MLD querier sends a group-specific query after receiving a leave report from a host. You can decrease this interval to reduce the amount of time it takes for multicast traffic to stop forwarding after the last member leaves a group.

You cannot configure this value directly for MLD snooping. MLD snooping inherits the value from the MLD value configured on the switch, which is applied to all VLANs on the switch.

To configure the MLD **query-last-member-interval**:

```
[edit protocols]
user@switch# set mld query-last-member-interval seconds
```

- **robust-count**—The number of times the querier resends a general membership query or a group-specific membership query (the default is 2 times). You can increase this count to tune for higher expected packet loss.

For MLD snooping, you can configure **robust-count** for a specific VLAN. If a VLAN does not have **robust-count** configured, the **robust-count** value is inherited from the value configured for MLD.

To configure **robust-count** for MLD snooping on a VLAN:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan-name robust-count number
```

The values configured for **query-interval**, **query-response-interval**, and **robust-count** determine the multicast listener interval—the length of time the switch waits for a group membership report after a general query before removing a multicast group from its multicast forwarding table. The switch calculates the multicast listener interval by multiplying **query-interval** by **robust-count** and then adding **query-response-interval**:

$(\text{query-interval} \times \text{robust-count}) + \text{query-response-interval} = \text{multicast listener interval}$

For example, the multicast listener interval is 260 seconds when the default settings for **query-interval**, **query-response-interval**, and **robust-count** are used:

$(125 \times 2) + 10 = 260$

You can display the time remaining in the multicast listener interval before a group times out by using the **show mld-snooping membership** command.

Related Documentation

- *Examples: Configuring MLD*

Configuring MLD Snooping Tracing Operations (CLI Procedure)

By enabling tracing operations for MLD snooping, you can record detailed messages about the operation of the protocol, such as the various types of protocol packets sent and received. [Table 5 on page 58](#) describes the tracing operations you can enable and the flags used to specify them in the tracing configuration.

Table 5: Supported Tracing Operations for MLD Snooping

| Tracing Operation | Flag |
|---|----------------|
| Trace all (equivalent of including all flags). | all |
| Trace general MLD snooping protocol events. | general |
| Trace communication over routing socket events. | krt |
| Trace leave reports. | leave |
| Trace next-hop-related events. | nexthop |
| Trace normal MLD snooping protocol events. If you do not specify this flag, only unusual or abnormal operations are traced. | normal |
| Trace all MLD packets. | packets |
| Trace policy processing. | policy |
| Trace MLD membership query messages. | query |
| Trace membership reports | report |
| Trace routing information. | route |
| Trace state transitions. | state |
| Trace routing protocol task processing. | task |
| Trace timer processing. | timer |
| Trace VLAN-related events. | vlan |

This topic covers:

- [Configuring Tracing Operations on page 59](#)
- [Viewing, Stopping, and Restarting Tracing Operations on page 59](#)

Configuring Tracing Operations

To configure tracing operations for MLD snooping:

1. Configure the filename for the trace file:

```
[edit protocols mld-snooping ]
user@switch# set traceoptions file filename
```

For example:

```
[edit protocols mld-snooping ]
user@switch# set traceoptions file mld-snoop-trace
```

2. (Optional) Configure the maximum number of trace files and size of the trace files:

```
[edit protocols mld-snooping ]
user@switch # set file files number size size
```

For example:

```
[edit protocols mld-snooping ]
user@switch # set traceoptions file files 5 size 1m
```

causes the contents of the trace file to be emptied and archived in a .gz file when the file reaches 1 MB. Four archive files are maintained, the contents of which are rotated whenever the current active trace file is archived.

If you omit this step, the maximum number of trace files defaults to 10, with the maximum file size defaulting to 128 K.

3. Specify one of the tracing flags shown in [Table 5 on page 58](#):

```
[edit protocols mld-snooping ]
user@switch # set traceoptions flag flagname
```

For example, to perform trace operations on VLAN-related events and MLD query messages:

```
[edit protocols mld-snooping ]
user@switch# set traceoptions flag vlan

[edit protocols mld-snooping ]
user@switch# set traceoptions flag query
```

Viewing, Stopping, and Restarting Tracing Operations

When you commit the configuration, tracing operations begin. You can view the trace file in the `/var/log` directory. For example:

```
user@switch> file show /var/log/mld-snoop-trace
```

You can stop and restart tracing operations by deactivating and reactivating the configuration:

```
[edit]
user@switch# deactivate protocols mld-snooping traceoptions
```

```
[edit]
user@switch# activate protocols mld-snooping traceoptions
```

- Related Documentation**
- [Configuring MLD Snooping on a VLAN \(CLI Procedure\) on page 50](#)
 - [Tracing and Logging Junos OS Operations](#)

Configuring Multicast VLAN Registration (CLI Procedure)

Multicast VLAN registration (MVR) enables hosts that are not part of a multicast source VLAN (MVLAN) to still receive multicast streams from the MVLAN, allowing an MVLAN to be shared across a Layer 2 network. Hosts remain in their own VLANs for bandwidth and security reasons but are able to receive multicast streams from the MVLAN.

You can configure one or more VLANs on a switch to be MVLANs or MVR receiver VLANs. By default, MVR is not configured on EX Series switches and the QFX Series.



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



NOTE: When you configure MVR, the following restrictions apply:

- You cannot enable multicast protocols on VLAN interfaces that are members of MVLANs.
- If you configure an MVLAN in proxy mode, IGMP snooping proxy mode is automatically enabled on all MVR receiver VLANs of this MVLAN. If a VLAN is an MVR receiver VLAN for multiple MVLANs, all of the MVLANs must have proxy mode enabled or all must have proxy mode disabled. You can enable proxy mode only on VLANs that are configured as MVR source VLANs and that are not configured for Q-in-Q tunneling.
- After you configure a VLAN as an MVLAN, that VLAN is no longer available for other uses.

To configure MVR:

1. Configure the VLAN named mv0 to be an MVLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan mv0 data-forwarding source groups
225.10.0.0/16
```

2. Configure the MVLAN mv0 to be a proxy VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan mv0 proxy source-address 10.0.0.1
```

3. Configure the VLAN named v2 to be an MVR receiver VLAN with mv0 as its source:

[edit protocols]

user@switch# set igmp-snooping vlan v2 data-forwarding receiver source-vlans mv0

4. Install forwarding entries in the MVR receiver VLAN:

[edit protocols]

user@switch# set igmp-snooping vlan mv0 data-forwarding receiver install

**Related
Documentation**

- [Example: Configuring Multicast VLAN Registration on page 33](#)
- [Understanding Multicast VLAN Registration on page 23](#)

CHAPTER 6

Configuration Statements

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 65](#)
- [\[edit protocols igmp\] Configuration Statement Hierarchy on EX Series Switches on page 67](#)
- [\[edit protocols igmp-snooping\] Configuration Statement Hierarchy on EX Series Switches on page 68](#)
- [accounting \(Protocols IGMP Interface\) on page 70](#)
- [accounting \(Protocols IGMP\) on page 70](#)
- [address \(Anycast RPs\) on page 71](#)
- [address \(Local RPs\) on page 72](#)
- [anycast-pim on page 73](#)
- [assert-timeout on page 74](#)
- [auto-rp on page 75](#)
- [bootstrap on page 76](#)
- [bootstrap-export on page 77](#)
- [bootstrap-import on page 78](#)
- [bootstrap-priority on page 79](#)
- [data-forwarding on page 80](#)
- [dense-groups on page 81](#)
- [disable \(IGMP Snooping\) on page 81](#)
- [disable \(Protocols IGMP\) on page 82](#)
- [disable \(MLD Snooping\) on page 82](#)
- [disable \(PIM\) on page 83](#)
- [dr-election-on-p2p on page 84](#)
- [dr-register-policy on page 84](#)
- [embedded-rp on page 85](#)
- [export \(Protocols PIM Bootstrap\) on page 86](#)
- [family \(Bootstrap\) on page 87](#)
- [family \(Local RP\) on page 88](#)

- [graceful-restart \(Protocols PIM\) on page 89](#)
- [group \(IGMP Snooping\) on page 89](#)
- [group \(Protocols IGMP\) on page 90](#)
- [group \(MLD Snooping\) on page 91](#)
- [group-ranges on page 92](#)
- [groups \(Multicast VLAN Registration\) on page 93](#)
- [hello-interval \(Protocols PIM\) on page 94](#)
- [hold-time \(Protocols PIM\) on page 95](#)
- [igmp-snooping on page 96](#)
- [immediate-leave \(Protocols IGMP\) on page 97](#)
- [immediate-leave \(IGMP Snooping\) on page 98](#)
- [immediate-leave \(MLD Snooping\) on page 99](#)
- [import \(Protocols PIM Bootstrap\) on page 100](#)
- [import \(Protocols PIM\) on page 101](#)
- [infinity on page 102](#)
- [install \(Multicast VLAN Registration\) on page 102](#)
- [interface \(IGMP Snooping\) on page 103](#)
- [interface \(Protocols PIM\) on page 104](#)
- [interface \(Protocols IGMP\) on page 106](#)
- [interface \(MLD Snooping\) on page 107](#)
- [join-load-balance on page 108](#)
- [local on page 109](#)
- [local-address \(Protocols PIM\) on page 110](#)
- [mapping-agent-election on page 111](#)
- [maximum-rps on page 112](#)
- [mld-snooping on page 113](#)
- [mode \(Protocols PIM\) on page 114](#)
- [multicast-router-interface \(IGMP Snooping\) on page 115](#)
- [multicast-router-interface \(MLD Snooping\) on page 116](#)
- [neighbor-policy on page 117](#)
- [pim on page 118](#)
- [priority \(PIM Interfaces\) on page 122](#)
- [priority \(Bootstrap\) on page 123](#)
- [priority \(PIM RPs\) on page 124](#)
- [promiscuous-mode \(Protocols IGMP\) on page 125](#)
- [proxy \(Multicast VLAN Registration\) on page 125](#)
- [query-interval \(Protocols IGMP\) on page 126](#)

- [query-last-member-interval \(Protocols IGMP\) on page 127](#)
- [query-response-interval \(Protocols IGMP\) on page 128](#)
- [receiver on page 129](#)
- [restart-duration \(Protocols PIM\) on page 130](#)
- [rib-group \(Protocols PIM\) on page 131](#)
- [robust-count \(IGMP Snooping\) on page 132](#)
- [robust-count \(Protocols IGMP\) on page 132](#)
- [robust-count \(MLD Snooping\) on page 133](#)
- [rp on page 134](#)
- [rp-register-policy on page 136](#)
- [rp-set on page 137](#)
- [source \(Multicast VLAN Registration\) on page 137](#)
- [source \(Protocols IGMP\) on page 138](#)
- [source-vlans on page 139](#)
- [spt-threshold on page 140](#)
- [ssm-map \(Protocols IGMP\) on page 141](#)
- [static \(IGMP Snooping\) on page 141](#)
- [static \(Protocols PIM\) on page 142](#)
- [static \(Protocols IGMP\) on page 143](#)
- [static \(MLD Snooping\) on page 144](#)
- [traceoptions \(Protocols PIM\) on page 145](#)
- [traceoptions \(Protocols IGMP\) on page 148](#)
- [traceoptions \(IGMP Snooping\) on page 151](#)
- [traceoptions \(MLD Snooping\) on page 153](#)
- [version \(Protocols IGMP\) on page 155](#)
- [version \(IGMP Snooping\) on page 156](#)
- [version \(MLD Snooping\) on page 157](#)
- [version \(PIM\) on page 158](#)
- [vlan \(IGMP Snooping\) on page 159](#)
- [vlan \(MLD Snooping\) on page 161](#)

[edit protocols] Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the **[edit protocols]** hierarchy:

- [\[edit protocols bfd\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols bgp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols connections\] Configuration Statement Hierarchy on EX Series Switches](#)

- [\[edit protocols dcbx\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols dot1x\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols igmp\]](#) Configuration Statement Hierarchy on EX Series Switches on page 67
- [\[edit protocols igmp-snooping\]](#) Configuration Statement Hierarchy on EX Series Switches on page 68
- [\[edit protocols isis\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols lacp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols link-management\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols lldp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols lldp-med\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols mld\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols mld-snooping\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols mpls\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols msdp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols mstp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols mvrp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols neighbor-discovery\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols oam\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols ospf\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols ospf3\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols pim\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols rip\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols ripng\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols router-advertisement\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols router-discovery\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols rstp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols rsvp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols sflow\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols stp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols uplink-failure-detection\]](#) Configuration Statement Hierarchy on EX Series Switches

- [\[edit protocols vrrp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols vstp\] Configuration Statement Hierarchy on EX Series Switches](#)

**Related
Documentation**

- [EX Series Switch Software Features Overview](#)
- [EX Series Virtual Chassis Software Features Overview](#)

[\[edit protocols igmp\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit protocols igmp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols igmp\] Hierarchy Level on page 67](#)
- [Unsupported Statements in the \[edit protocols igmp\] Hierarchy Level on page 68](#)

Supported Statements in the **[edit protocols igmp]** Hierarchy Level

The following hierarchy shows the **[edit protocols igmp]** configuration statements supported on EX Series switches:

```
protocols {
  igmp {
    accounting;
    interface interface-name {
      (accounting | no-accounting);
      disable;
      group-policy [ policy-names ];
      group-policy policy-name;
      immediate-leave;
      oif-map [ map-names ];
      passive <allow-receive> <send-general-query> <send-group-query>;
      promiscuous-mode;
      ssm-map ssm-map-name;
      static {
        group multicast-group-address {
          exclude;
          group-count number;
          group-increment increment;
          source ip-address {
            source-count number;
            source-increment increment;
          }
        }
      }
    }
  }
}
```

```
    }  
  }  
}  
version version;  
}  
maximum-transmit-rate packets-per-second;  
query-interval seconds;  
query-last-member-interval seconds;  
query-response-interval seconds;  
robust-count number;  
traceoptions {  
  file filename <files number> <size maximum-file-size> <world-readable |  
    no-world-readable>;  
  flag flag <flag-modifier> <disable>;  
}  
}
```

Unsupported Statements in the [edit protocols igmp] Hierarchy Level

All statements in the [edit protocols igmp] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 65](#)

[edit protocols igmp-snooping] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols igmp-snooping] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols igmp-snooping\] Hierarchy Level on page 68](#)
- [Unsupported Statements in the \[edit protocols igmp-snooping\] Hierarchy Level on page 69](#)

Supported Statements in the [edit protocols igmp-snooping] Hierarchy Level

The following hierarchy shows the [edit protocols igmp-snooping] configuration statements supported on EX Series switches:

```

protocols {
  igmp-snooping {
    vlan vlan-identifier {
      immediate-leave;
      interface (all | interface-name) {
        group-limit <1..65535>
        host-only-interface
        multicast-router-interface;
        immediate-leave;
        static {
          group multicast-ip-address; {
            source <>
          }
        }
      }
    }
  }
  proxy {
    source-address ip-address;
  }
  query-interval number;
  query-last-member-interval number;
  query-response-interval number;
  robust-count number;
  traceoptions {
    file filename <files number> <no-stamp> <replace> <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}

```

Unsupported Statements in the [edit protocols igmp-snooping] Hierarchy Level

All statements in the [edit protocols igmp-snooping] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 65](#)

accounting (Protocols IGMP Interface)

| | |
|---------------------------------|--|
| Syntax | (accounting no-accounting); |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp interface <i>interface-name</i>], [edit protocols igmp interface <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Enable or disable the collection of IGMP join and leave event statistics for an interface. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Recording IGMP Join and Leave Events</i> |

accounting (Protocols IGMP)

| | |
|---------------------------------|--|
| Syntax | accounting; |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp], [edit protocols igmp] |
| Release Information | Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Enable the collection of IGMP join and leave event statistics on the system. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Recording IGMP Join and Leave Events</i> |

address (Anycast RPs)

| | |
|---------------------------------|--|
| Syntax | <code>address <i>address</i> <forward-msdp-sa>;</code> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols <code>pim rp local</code> (inet inet6) <code>anycast-pim rp-set</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim rp local</code> (inet inet6) <code>anycast-pim rp-set</code>],</p> <p>[edit protocols <code>pim rp local</code> (inet inet6) <code>anycast-pim rp-set</code>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim rp local</code> (inet inet6) <code>anycast-pim rp-set</code>]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Configure the anycast rendezvous point (RP) addresses in the RP set. Multiple addresses can be configured in an RP set. If the RP has peer Multicast Source Discovery Protocol (MSDP) connections, then the RP must forward MSDP source active (SA) messages. |
| Options | <p><i>address</i>—RP address in an RP set.</p> <p><i>forward-msdp-sa</i>—(Optional) Forward MSDP SAs to this address.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |

address (Local RPs)

| | |
|---------------------------------|--|
| Syntax | <code>address <i>address</i>;</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols <code>pim rp local family</code> (inet inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim rp local family</code> (inet inet6)], [edit protocols <code>pim rp local family</code> (inet inet6)], [edit routing-instances <i>routing-instance-name</i> protocols <code>pim rp local family</code> (inet inet6)] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Configure the local rendezvous point (RP) address. |
| Options | <i>address</i> —Local RP address. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring Local PIM RPs</i> |

anycast-pim

| | |
|---------------------------------|--|
| Syntax | <pre>anycast-pim { rp-set { address address <forward-msdp-sa>; } }</pre> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp local family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6)],</p> <p>[edit protocols pim rp local family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6)]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Configure properties for anycast RP using PIM.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Example: Configuring PIM Anycast With or Without MSDP</i> |

assert-timeout

| | |
|---------------------------------|--|
| Syntax | <code>assert-timeout <i>seconds</i>;</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Multicast routing devices running PIM sparse mode often forward the same stream of multicast packets onto the same LAN through the rendezvous-point tree (RPT) and shortest-path tree (SPT). PIM assert messages help routing devices determine which routing device forwards the traffic and prunes the RPT for this group. By default, routing devices enter an assert cycle every 180 seconds. You can configure this assert timeout to be between 5 and 210 seconds. |
| Options | <i>seconds</i> —Time for routing device to wait before another assert message cycle. Range: 5 through 210 seconds Default: 180 seconds |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring the PIM Assert Timeout</i> |

auto-rp

| | |
|---------------------------------|--|
| Syntax | <pre>auto-rp { (announce discovery mapping); (mapping-agent-election no-mapping-agent-election); }</pre> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp],</p> <p>[edit protocols pim rp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 7.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Configure automatic RP announcement and discovery. |
| Options | <p>announce—Configure the routing device to listen only for mapping packets and also to advertise itself if it is an RP.</p> <p>discovery—Configure the routing device to listen only for mapping packets.</p> <p>mapping—Configures the routing device to announce, listen for and generate mapping packets, and announce that the routing device is eligible to be an RP.</p> <p>The remaining statement is explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring PIM Auto-RP</i> |

bootstrap

| | |
|---------------------------------|---|
| Syntax | <pre>bootstrap { family (inet inet6) { export [<i>policy-names</i>]; import [<i>policy-names</i>]; priority <i>number</i>; } }</pre> |
| Hierarchy Level | <pre>[edit logical-systems <i>logical-system-name</i> protocols pim rp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp], [edit protocols pim rp], [edit routing-instances <i>routing-instance-name</i> protocols pim rp]</pre> |
| Release Information | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Configure parameters to control bootstrap routers and messages.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring PIM Bootstrap Properties for IPv4</i>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i> |

bootstrap-export

| | |
|---------------------------------|--|
| Syntax | <code>bootstrap-export [<i>policy-names</i>];</code> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp],</p> <p>[edit protocols pim rp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Apply one or more export policies to control outgoing PIM bootstrap messages. |
| Options | <i>policy-names</i> —Name of one or more import policies. |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring PIM Bootstrap Properties for IPv4</i> • <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i> • bootstrap-import on page 78 |

bootstrap-import

| | |
|---------------------------------|---|
| Syntax | <code>bootstrap-import [<i>policy-names</i>];</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim rp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp], [edit protocols pim rp], [edit routing-instances <i>routing-instance-name</i> protocols pim rp] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Apply one or more import policies to control incoming PIM bootstrap messages. |
| Options | <i>policy-names</i> —Name of one or more import policies. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring PIM Bootstrap Properties for IPv4</i>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i>• bootstrap-export on page 77 |

bootstrap-priority

| | |
|---------------------------------|---|
| Syntax | <code>bootstrap-priority <i>number</i>;</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim rp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp], [edit protocols pim rp], [edit routing-instances <i>routing-instance-name</i> protocols pim rp] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Configure whether this routing device is eligible to be a bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router. |
| Options | <i>number</i> —Priority for becoming the bootstrap router. A value of 0 means that the routing device is not eligible to be the bootstrap router. Range: 0 through 255 Default: 0 |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring PIM Bootstrap Properties for IPv4</i> |

data-forwarding

| | |
|---------------------------------|--|
| Syntax | <pre>data-forwarding { receiver { source-vlans <i>vlan-list</i>; install; } source { groups <i>group-prefix</i>; } }</pre> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series. |
| Description | <p>Configure the VLAN to be a multicast source VLAN (MVLAN) or a multicast VLAN registration (MVR) receiver VLAN. Each data-forwarding VLAN, which can be a multicast source VLAN (MVLAN) or a multicast receiver VLAN, must have exactly one source statement or exactly one receiver statement. A data-forwarding VLAN can operate only in IGMP version 2 (IGMPv2) mode.</p> <p>The remaining statements are explained separately.</p> |
| Default | Disabled |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring Multicast VLAN Registration on page 33• Configuring Multicast VLAN Registration (CLI Procedure) on page 60 |

dense-groups

| | |
|---------------------------------|---|
| Syntax | <code>dense-groups { addresses; }</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Configure which groups are operating in dense mode. |
| Options | addresses —Address of groups operating in dense mode. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring PIM Sparse-Dense Mode Properties |

disable (IGMP Snooping)

| | |
|---------------------------------|--|
| Syntax | <code>disable;</code> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 9.2 for EX Series switches. |
| Description | Disable IGMP snooping on the VLAN. Multicast traffic will be flooded to all interfaces on the VLAN except the source interface. |
| Default | If you do not include this statement in the configuration for a VLAN, IGMP snooping is enabled on the VLAN. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IGMP Snooping (CLI Procedure) on page 39 • show igmp-snooping vlans on page 227 |

disable (Protocols IGMP)

| | |
|---------------------------------|--|
| Syntax | disable; |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp interface interface-name], [edit protocols igmp interface interface-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Disable IGMP on the system. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Disabling IGMP |

disable (MLD Snooping)

| | |
|---------------------------------|---|
| Syntax | disable; |
| Hierarchy Level | [edit protocols mld-snooping vlan (all <i>vlan-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 12.1 for EX Series switches. |
| Description | Disable MLD snooping on the VLAN. Multicast traffic will be flooded to all interfaces in the VLAN except the source interface. |
| Default | If you do not include this statement, MLD snooping is enabled on all interfaces in the VLAN. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50• show mld-snooping vlans on page 238 |

disable (PIM)

| | |
|---------------------------------|--|
| Syntax | disable; |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp local family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6)],</p> <p>[edit protocols pim],</p> <p>[edit protocols pim family (inet inet6)],</p> <p>[edit protocols pim interface <i>interface-name</i>],</p> <p>[edit protocols pim rp local family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6)]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>disable statement extended to the [family] hierarchy level in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Explicitly disable PIM at the protocol, interface or family hierarchy levels. |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Disabling PIM disable (PIM Graceful Restart) |

dr-election-on-p2p

| | |
|---------------------------------|---|
| Syntax | dr-election-on-p2p; |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim] |
| Release Information | Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Enable PIM designated router (DR) election on point-to-point (P2P) links. |
| Default | No PIM DR election is performed on point-to-point links. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring PIM Designated Router Election on Point-to-Point Links</i> |

dr-register-policy

| | |
|---------------------------------|---|
| Syntax | dr-register-policy [<i>policy-names</i>]; |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim rp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp], [edit protocols pim rp], [edit routing-instances <i>routing-instance-name</i> protocols pim rp] |
| Release Information | Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Apply one or more policies to control outgoing PIM register messages. |
| Options | <i>policy-names</i> —Name of one or more import policies. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring Register Message Filters on a PIM RP and DR</i>• rp-register-policy on page 136 |

embedded-rp

| | |
|---------------------------------|--|
| Syntax | <pre> embedded-rp { group-ranges { destination-ip-prefix</prefix-length>; } maximum-rps limit; } </pre> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp],</p> <p>[edit protocols pim rp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Configure properties for embedded IP version 6 (IPv6) RPs.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring PIM Embedded RP for IPv6</i> |

export (Protocols PIM Bootstrap)

| | |
|---------------------------------|---|
| Syntax | <code>export [<i>policy-names</i>];</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim rp bootstrap family (inet inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap family (inet inet6)], [edit protocols pim rp bootstrap family (inet inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap family (inet inet6)] |
| Release Information | Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Apply one or more export policies to control outgoing PIM bootstrap messages. |
| Options | <i>policy-names</i> —Name of one or more import policies. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring PIM Bootstrap Properties for IPv4</i>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i>• import (Protocols PIM Bootstrap) on page 100 |

family (Bootstrap)

| | |
|---------------------------------|--|
| Syntax | <pre>family (inet inet6) { export [<i>policy-names</i>]; import [<i>policy-names</i>]; priority <i>number</i>; }</pre> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bootstrap],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap],</p> <p>[edit protocols pim rp bootstrap],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Configure which IP protocol type bootstrap properties to apply. |
| Options | <p>inet—Apply IP version 4 (IPv4) local RP properties.</p> <p>inet6—Apply IPv6 local RP properties.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring PIM Bootstrap Properties for IPv4</i> • <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i> |

family (Local RP)

| | |
|---------------------------------|--|
| Syntax | <pre> family (inet inet6) { disable; address address; anycast-pim { local-address address; rp-set { address address <forward-msdp-sa>; } } group-ranges { destination-ip-prefix </prefix-length>; } hold-time seconds; override; priority number; } </pre> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp local],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp local],</p> <p>[edit protocols pim rp local],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp local]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Configure which IP protocol type local RP properties to apply. |
| Options | <p>inet—Apply IP version 4 (IPv4) local RP properties.</p> <p>inet6—Apply IPv6 local RP properties.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Configuring Local PIM RPs |


graceful-restart (Protocols PIM)

| | |
|---------------------------------|---|
| Syntax | graceful-restart { disable; no-bidirectional-mode; restart-duration seconds; } |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Configure PIM sparse mode graceful restart. The remaining statements are explained separately. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring PIM Sparse Mode Graceful Restart |

group (IGMP Snooping)

| | |
|---------------------------------|---|
| Syntax | group <i>ip-address</i> ; |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>) interface (all <i>interface-name</i>) static] |
| Release Information | Statement introduced in Junos OS Release 9.1 for EX Series switches. |
| Description | Configure a static multicast group on an interface. |
| Options | <i>ip-address</i> —Valid IP multicast address for the multicast group. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IGMP Snooping (CLI Procedure) on page 39 • show igmp-snooping membership on page 220 |

group (Protocols IGMP)

| | |
|---|---|
| Syntax | <pre>group <i>multicast-group-address</i> { exclude; group-count <i>number</i>; group-increment <i>increment</i>; source <i>ip-address</i> { source-count <i>number</i>; source-increment <i>increment</i>; } }</pre> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp interface <i>interface-name</i> static], [edit protocols igmp interface <i>interface-name</i> static] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify the IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface. |
| <hr/> | |
| <div> NOTE: You must specify a unique address for each group.</div> <hr/> | |
| The remaining statements are explained separately. | |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Enabling IGMP Static Group Membership</i> |

group (MLD Snooping)

| | |
|---------------------------------|--|
| Syntax | <code>group <i>multicast-group-address</i> { source <i>ip-address</i>; }</code> |
| Hierarchy Level | [edit protocols mld-snooping vlan (all <i>vlan-name</i>) interface (all <i>interface-name</i>) static] [edit routing-instances <i>instance-name</i> protocols mld-snooping vlan <i>vlan-name</i> interface <i>interface-name</i> static] |
| Release Information | Statement introduced in Junos OS Release 12.1 for EX Series switches. Support at the [edit routing-instances <i>instance-name</i> protocols mld-snooping vlan <i>vlan-name</i> interface <i>interface-name</i> static] hierarchy level introduced in Junos OS Release 13.3 for EX Series switches. Support for the source statement introduced in Junos OS Release 13.3 for EX Series switches. |
| Description | Configure a static multicast group on an interface and (optionally) the source address for the multicast group. |
| Options | <i>multicast-group-address</i> —Valid IP multicast address for the multicast group. <i>source ip-address</i> —Valid IP multicast address for the source of the multicast group. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |

group-ranges

| | |
|---------------------------------|--|
| Syntax | <pre>group-ranges { destination-ip-prefix</prefix-length>; }</pre> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp embedded-rp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp embedded-rp],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim rp embedded-rp],</p> <p>[edit protocols pim rp local family (inet inet6)],</p> <p>[edit protocols pim rp static address <i>address</i>],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp embedded-rp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp static address <i>address</i>]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 13.3 for the PTX5000 router.</p> |
| Description | Configure the address ranges of the multicast groups for which this routing device can be a rendezvous point (RP). |
| Default | The routing device is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). |
| Options | <i>destination-ip-prefix</prefix-length></i> —Addresses or address ranges for which this routing device can be an RP. |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring Local PIM RPs</i> • <i>Configuring PIM Embedded RP for IPv6</i> • <i>Example: Configuring Bidirectional PIM</i> |

groups (Multicast VLAN Registration)

| | |
|---------------------------------|--|
| Syntax | <code>groups <i>group-prefix</i>;</code> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>) data-forwarding source] |
| Release Information | Statement introduced in Junos OS Release 9.6 for EX Series switches. |
| Description | Specify the IP address range of the multicast VLAN (MVLAN) source interfaces. |
| Default | Disabled |
| Options | <i>group-prefix</i> —IP address range of the source group. Each MVLAN must have exactly one groups statement. If there are multiple MVLANs on the switch, their group ranges must be unique. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring Multicast VLAN Registration on page 33• Configuring Multicast VLAN Registration (CLI Procedure) on page 60 |

hello-interval (Protocols PIM)

| | |
|---------------------------------|---|
| Syntax | <code>hello-interval seconds;</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface interface-name], [edit protocols pim interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols pim interface interface-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify how often the routing device sends PIM hello packets out of an interface. |
| Options | seconds —Length of time between PIM hello packets. Range: 0 through 255 Default: 30 seconds |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• hold-time on page 95• <i>Modifying the PIM Hello Interval</i> |

hold-time (Protocols PIM)

| | |
|---------------------------------|---|
| Syntax | <code>hold-time seconds;</code> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim rp local family (inet inet6)],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6)]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 13.3 for the PTX5000 router.</p> |
| Description | Specify the time period for which a neighbor is to consider the sending routing device (this routing device) to be operative (up). |
| Options | <p>seconds—Hold time.</p> <p>Range: 0 through 255</p> <p>Default: 150 seconds</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring Local PIM RPs in the Multicast Protocols Feature Guide for Routing Devices</i> • <i>Example: Configuring Bidirectional PIM</i> |

igmp-snooping

```
Syntax  igmp-snooping {
        traceoptions {
            file filename <files number> <no-stamp> <replace> <size size> <world-readable |
            no-world-readable>;
            flag flag <flag-modifier>;
        }
        vlan (all | vlan-name) {
            data-forwarding {
                source {
                    groups group-prefix;
                }
                receiver {
                    source-vlans vlan-list;
                    install;
                }
            }
            disable;
            immediate-leave;
            interface (all | interface-name) {
                multicast-router-interface;
                static {
                    group ip-address;
                }
            }
            proxy {
                source-address ip-address;
            }
            robust-count number;
            version version;
        }
    }
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure IGMP snooping. The factory default configuration enables IGMP snooping on all VLANs.


The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation


- [Example: Configuring IGMP Snooping on EX Series Switches on page 27](#)
- [Configuring IGMP Snooping \(CLI Procedure\) on page 39](#)

immediate-leave (Protocols IGMP)


| | |
|---------------------------------|--|
| Syntax | <code>immediate-leave;</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp interface <i>interface-name</i>], [edit protocols igmp interface <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | <p>The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.</p> <p>Starting in Junos OS Release 9.3, both IGMP version 2 and IGMP version 3 do host tracking when the immediate-leave statement is configured. This means that the multicast group leaves only when the last host leaves. The routing device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</p> <p>When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.</p> <p>When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.</p> |
| | <p> NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the routing device only knows about the one interested host and does not know about the others.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |

- Related Documentation**
- [Specifying Immediate-Leave Host Removal for IGMP](#)

immediate-leave (IGMP Snooping)

| | |
|---------------------------------|---|
| Syntax | <code>immediate-leave;</code> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 9.1 for EX Series switches. |
| Description | <p>Configure IGMP snooping immediate leave for the specified VLAN. When you configure immediate leave, host tracking is enabled, which allows the switch to track the hosts that send membership reports. The switch can then determine when the last host on an interface leaves the multicast group and immediately stop forwarding multicast traffic to the interface.</p> <p>Configuring immediate leave reduces the amount of time it takes for the switch to stop sending multicast traffic to an interface when the last host leaves the group. When immediate leave is disabled, the switch no longer tracks hosts. Instead, whenever it receives a leave report from a host, it sends out a group-specific query to all hosts. If it does not receive any membership reports on the interface in response to the group-specific query within a set interval, it stops forwarding multicast traffic to the interface.</p> |
| | <p> NOTE: Immediate leave is supported for both IGMP version 2 (IGMPv2) and IGMPv3. However, with IGMPv2, we recommend that you configure immediate leave only when there is only one IGMP host on an interface. In IGMPv2, only one host on a interface sends a membership report in response to a general query—any other interested hosts suppress their reports. Report suppression avoids a flood of reports for the same group, but it also interferes with host tracking because the switch knows only about one interested host on the interface at any given time.</p> |
| Default | The immediate-leave feature is disabled. |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Example: Configuring IGMP Snooping on EX Series Switches on page 27 • Configuring IGMP Snooping (CLI Procedure) on page 39 |

immediate-leave (MLD Snooping)

| | |
|---------------------------------|---|
| Syntax | <code>immediate-leave;</code> |
| Hierarchy Level | <code>[edit protocols mld-snooping vlan (all <i>vlan-name</i>)]</code> <code>[edit protocols mld-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>]</code> <code>[edit routing-instances <i>instance-name</i> protocols mld-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>]</code> |
| Release Information | <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Support at the <code>[edit protocols mld-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>]</code> and the <code>[edit routing-instances <i>instance-name</i> protocols mld-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>]</code> hierarchy levels introduced in Junos OS Release 13.3 for EX Series switches.</p> <p>Support at</p> |
| Description | <p>Configure MLD snooping immediate leave for the specified VLAN or interface. When you configure immediate leave, host tracking is enabled, which allows the switch to track the hosts that send join messages. The switch can then determine when the last host on an interface leaves the multicast group and immediately stop forwarding multicast traffic to the interface.</p> <p>Configuring immediate leave reduces the amount of time it takes for the switch to stop sending multicast traffic to an interface when the last host leaves the group. When immediate leave is disabled, the switch no longer tracks hosts. Instead, whenever it receives a leave message from a host, it sends out a group membership query to all hosts. If it does not receive any join group reports on the interface in response to the group membership query within a set interval, it then stops forwarding multicast traffic to the interface.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 20px;"> <p> NOTE: Immediate leave is supported for both MLD version 1 (MLDv1) and MLDv2. However, with MLDv1, we recommend that you configure immediate leave only when there is only one MLD host on an interface. In MLDv1, only one host on a interface sends a join report in response to a group membership query—any other interested hosts suppress their reports. This report-suppression feature means that the switch only knows about one interested host on the interface at any given time.</p> </div> |
| Default | The immediate-leave feature is disabled. |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |

import (Protocols PIM Bootstrap)

| | |
|---------------------------------|---|
| Syntax | <code>import [<i>policy-names</i>];</code> |
| Hierarchy Level | <code>[edit logical-systems <i>logical-system-name</i> protocols pim rp bootstrap (inet inet6)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> pim rp bootstrap (inet inet6)],</code> <code>[edit protocols pim rp bootstrap (inet inet6)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap (inet inet6)]</code> |
| Release Information | Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Apply one or more import policies to control incoming PIM bootstrap messages. |
| Options | <i>policy-names</i> —Name of one or more import policies. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring PIM Bootstrap Properties for IPv4</i>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i>• export (Protocols PIM Bootstrap) on page 86 |

import (Protocols PIM)

| | |
|---------------------------------|--|
| Syntax | <code>import [<i>policy-names</i>];</code> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],</p> <p>[edit protocols pim],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Apply one or more policies to routes being imported into the routing table from PIM. Use the import statement to filter PIM join messages and prevent them from entering the network. |
| Options | <i>policy-names</i> —Name of one or more policies. |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Filtering Incoming PIM Join Messages</i> |

infinity

| | |
|---------------------------------|---|
| Syntax | <code>infinity [<i>policy-names</i>];</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim spt-threshold], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim spt-threshold], [edit protocols pim spt-threshold], [edit routing-instances <i>routing-instance-name</i> protocols pim spt-threshold] |
| Release Information | Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the infinity statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair. |
| Options | <i>policy-names</i> —Name of one or more policies. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring the PIM SPT Threshold Policy |

install (Multicast VLAN Registration)

| | |
|---------------------------------|--|
| Syntax | <code>install;</code> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>) data-forwarding receiver] |
| Release Information | Statement introduced in Junos OS Release 9.6 for EX Series switches. |
| Description | Install forwarding entries in the multicast receiver VLAN. By default, the multicast VLAN (MVLAN) installs forwarding entries for MVLAN groups only. |
| Default | Disabled |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring Multicast VLAN Registration on page 33• Configuring Multicast VLAN Registration (CLI Procedure) on page 60 |

interface (IGMP Snooping)

| | |
|---------------------------------|---|
| Syntax | <pre> interface (all <i>interface-name</i>) { multicast-router-interface; static { group <i>ip-address</i>; } } </pre> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 9.1 for EX Series switches. |
| Description | For IGMP snooping, configure an interface as either a multicast-router interface or as a static member of a multicast group. |
| Options | <p>all—All interfaces in the VLAN.</p> <p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Example: Configuring IGMP Snooping on EX Series Switches on page 27 • Configuring IGMP Snooping (CLI Procedure) on page 39 • show igmp-snooping vlans on page 227 |

interface (Protocols PIM)

```
Syntax interface (Protocols PIM) (all | interface-name) {
    accept-remote-source;
    disable;
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (0 | 1 | automatic);
    }
    bidirectional {
        df-election {
            backoff-period milliseconds;
            offer-period milliseconds;
            robustness-count number;
        }
    }
    family (inet | inet6) {
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (0 | 1 | automatic);
        }
        disable;
    }
    hello-interval seconds;
}
```

```

mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse | sparse-dense);
neighbor-policy [ policy-names ];
override-interval milliseconds;
priority number;
propagation-delay milliseconds;
reset-tracking-bit;
version version;
}

```

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim**],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
pim],
 [edit protocols **pim**],
 [edit routing-instances *routing-instance-name* protocols **pim**]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Enable PIM on an interface and configure interface-specific properties.

Options *interface-name*—Name of the interface. Specify the full interface name, including the
 physical and logical address components. To configure all interfaces, you can specify
all.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation • *PIM on Aggregated Interfaces*

interface (Protocols IGMP)

| | |
|---------------------------------|--|
| Syntax | <pre> interface <i>interface-name</i> { disable; (accounting no-accounting); group-limit <i>limit</i>; group-policy [<i>policy-names</i>]; immediate-leave; oif-map <i>map-name</i>; passive; promiscuous-mode; ssm-map <i>ssm-map-name</i>; ssm-map-policy <i>ssm-map-policy-name</i>; static { group <i>multicast-group-address</i> { exclude; group-count <i>number</i>; group-increment <i>increment</i>; source <i>ip-address</i> { source-count <i>number</i>; source-increment <i>increment</i>; } } } version <i>version</i>; } </pre> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp], [edit protocols igmp] |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Enable IGMP on an interface and configure interface-specific properties. |
| Options | <p><i>interface-name</i>—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify all.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Enabling IGMP |

interface (MLD Snooping)

| | |
|---------------------------------|---|
| Syntax | <pre> interface (all <i>interface-name</i>) { group-limit <i>limit</i>; host-only-interface; immediate-leave; multicast-router-interface; static { group <i>ip-address</i> { source <i>ip-address</i>; } } } </pre> |
| Hierarchy Level | [edit protocols mld-snooping <i>vlan</i> (all <i>vlan-name</i>)] [edit routing-instances <i>instance-name</i> protocols mld-snooping <i>vlan</i> (<i>vlan-name</i>)] |
| Release Information | <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Support at the [edit routing-instances <i>instance-name</i> protocols mld-snooping <i>vlan</i> <i>vlan-name</i>] hierarchy introduced in Junos OS Release 13.3 for EX Series switches.</p> <p>Support for the group-limit, host-only-interface, and the immediate-leave statements introduced in Junos OS Release 13.3 for EX Series switches.</p> |
| Description | For MLD snooping, configure an interface as a static multicast-router interface, a host-side interface, or a static member of a multicast group. |
| Options | <p>all—(All EX Series switches except EX9200) All interfaces in the VLAN.</p> <p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |

join-load-balance

| | |
|---------------------------------|---|
| Syntax | <pre>join-load-balance { automatic; }</pre> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim] |
| Release Information | Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Enable load balancing of PIM join messages across interfaces and routing devices. |
| Options | automatic —Enables automatic load balancing of PIM join messages. When a new interface or neighbor is introduced into the network, ECMP joins are redistributed with minimal disruption to traffic. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring PIM Make-Before-Break Join Load Balancing</i>• <i>Configuring PIM Join Load Balancing</i>• <i>clear pim join-distribution</i> in the CLI Explorer |

local

| | |
|---------------------------------|---|
| Syntax | <pre> local { disable; address address; family (inet inet6) { disable; address address; anycast-pim { local-address address; rp-set { address address <forward-msdp-sa>; } } group-ranges { destination-ip-prefix</prefix-length>; } hold-time seconds; override; priority number; } group-ranges { destination-ip-prefix</prefix-length>; } hold-time seconds; override; priority number; } </pre> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp],</p> <p>[edit protocols pim rp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Configure the routing device's RP properties.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring Local PIM RPs</i> |

local-address (Protocols PIM)

| | |
|---------------------------------|---|
| Syntax | <code>local-address address;</code> |
| Hierarchy Level | <code>[edit logical-systems <i>logical-system-name</i> protocols pim rp local family (inet inet6) anycast-pim],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6) anycast-pim],</code> <code>[edit protocols pim rp local family (inet inet6) anycast-pim],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6) anycast-pim]</code> |
| Release Information | Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Configure the routing device local address for the anycast rendezvous point (RP). If this statement is omitted, the router ID is used as this address. |
| Options | address —Anycast RP IPv4 or IPv6 address, depending on family configuration. |
| Required Privilege Level | <code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring PIM Anycast With or Without MSDP</i> |

mapping-agent-election

| | |
|---------------------------------|---|
| Syntax | (mapping-agent-election no-mapping-agent-election); |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim rp auto-rp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp auto-rp], [edit protocols pim rp auto-rp], [edit routing-instances <i>routing-instance-name</i> protocols pim rp auto-rp] |
| Release Information | Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Configure the routing device mapping announcements as a mapping agent. |
| Options | mapping-agent-election —Mapping agents do not announce mappings when receiving mapping messages from a higher-addressed mapping agent. no-mapping-agent-election —Mapping agents always announce mappings and do not perform mapping agent election. Default: mapping-agent-election |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring PIM Auto-RP</i> |

maximum-rps

| | |
|---------------------------------|---|
| Syntax | <code>maximum-rps <i>limit</i>;</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim rp embedded-rp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp embedded-rp], [edit protocols pim rp embedded-rp], [edit routing-instances <i>routing-instance-name</i> protocols pim rp embedded-rp] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Limit the number of RPs that the routing device acknowledges. |
| Options | <i>limit</i> —Number of RPs. Range: 1 through 500 Default: 100 |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring PIM Embedded RP for IPv6</i> |


mld-snooping

| | |
|---------------------------------|--|
| Syntax | <pre> mld-snooping { traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <replace> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>>; } vlan (all <i>vlan-name</i>) { disable; immediate-leave; interface (all <i>interface-name</i>) { multicast-router-interface; static { group <i>ip-address</i>; } } robust-count <i>number</i>; version <i>version</i>; } } </pre> |
| Hierarchy Level | [edit protocols] |
| Release Information | Statement introduced in Junos OS Release 12.1 for EX Series switches. |
| Description | <p>Enable and configure MLD snooping.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Example: Configuring MLD Snooping on page 30 • Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |


mode (Protocols PIM)

| | |
|---------------------------------|---|
| Syntax | <code>mode (bidirectional-sparse bidirectional-sparse-dense dense sparse sparse-dense);</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>], [edit protocols pim interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. bidirectional-sparse and bidirectional-sparse-dense options introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 13.3 for the PTX5000 router. |
| Description | Configure the PIM mode on the interface. |
| Options | <p>The choice of PIM mode is closely tied to controlling how groups are mapped to PIM modes, as follows:</p> <ul style="list-style-type: none">• bidirectional-sparse—Use if all multicast groups are operating in bidirectional, sparse, or SSM mode.• bidirectional-sparse-dense—Use if multicast groups, except those that are specified in the dense-groups statement, are operating in bidirectional, sparse, or SSM mode.• dense—Use if all multicast groups are operating in dense mode.• sparse—Use if all multicast groups are operating in sparse mode or SSM mode.• sparse-dense—Use if multicast groups, except those that are specified in the dense-groups statement, are operating in sparse mode or SSM mode. <p>Default: Sparse mode</p> |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring PIM Dense Mode Properties in the Multicast Protocols Feature Guide for Routing Devices</i>• <i>Configuring PIM Sparse-Dense Mode Properties in the Multicast Protocols Feature Guide for Routing Devices</i>• <i>Example: Configuring Bidirectional PIM</i> |

multicast-router-interface (IGMP Snooping)

| | |
|--|---|
| Syntax | multicast-router-interface; |
| Hierarchy Level | <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols igmp-snooping vlan (all <i>vlan-name</i>) interface (all <i>interface-name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping interface <i>interface-name</i>]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p> |
| Description | <p>Statically configure the interface as an IGMP snooping multicast-router interface—that is, an interface that faces toward a multicast router or other IGMP querier.</p> |
| <div>  <p>NOTE: If the specified interface is a trunk port, the interface becomes a multicast-routing device interface for all VLANs configured on the trunk port. In addition, all unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded to the multicast routing device interface, even if the interface is configured as a multicast routing device interface only for IGMP snooping.</p> <p>Configure an interface as a bridge interface toward other multicast routing devices.</p> </div> | |
| Default | The interface can either be a host-side or multicast-routing device interface. |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Example: Configuring IGMP Snooping</i> • <i>IGMP Snooping in MC-LAG Active-Active Mode on MX Series Routers Overview</i> • <i>host-only-interface</i> |

multicast-router-interface (MLD Snooping)

| | |
|--|---|
| Syntax | multicast-router-interface; |
| Hierarchy Level | [edit protocols mld-snooping vlan (all <i>vlan-name</i>) interface (all <i>interface-name</i>)] [edit routing-instances <i>instance-name</i> protocols mld-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.1 for EX Series switches. Support at the [edit routing-instances <i>instance-name</i> protocols mld-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>] hierarchy level introduced in Junos OS Release 13.3 for EX Series switches. |
| Description | Statically configure the interface as a multicast-router interface—that is, an interface that faces towards a multicast router or other MLD querier. |
| <div> NOTE: If the specified interface is a trunk port, the interface becomes a multicast-router interface for all VLANs configured on the trunk port. In addition, all unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded to the multicast router interface, even if the interface is configured as a multicast-router interface only for MLD snooping.</div> | |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |

neighbor-policy

| | |
|---------------------------------|---|
| Syntax | <code>neighbor-policy [<i>policy-names</i>];</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface interface-name], [edit protocols pim interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols pim interface interface-name] |
| Release Information | Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Apply a PIM interface-level policy to filter neighbor IP addresses. |
| Options | <i>policy-name</i> —Name of the policy that filters neighbor IP addresses. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring Interface-Level PIM Neighbor Policies</i> |

pim

```
Syntax  pim {
    disable;
    assert-timeout seconds;
    dense-groups {
        addresses;
    }
    dr-election-on-p2p;
    export;
    family (inet | inet6) {
        disable;
    }
    graceful-restart {
        disable;
        no-bidirectional-mode;
        restart-duration seconds;
    }
    import [ policy-names ];
    interface interface-name {
        family (inet | inet6) {
            disable;
        }
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
            }
            loose-check;
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (0 | 1 | automatic);
        }
        accept-remote-source;
        disable;
        bidirectional {
            df-election {
                backoff-period milliseconds;
                offer-period milliseconds;
                robustness-count number;
            }
        }
        family (inet | inet6) {
            disable;
        }
        hello-interval seconds;
    }
}
```

```

mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse |
sparse-dense);
neighbor-policy [ policy-names ];
override-interval milliseconds;
priority number;
propagation-delay milliseconds;
reset-tracking-bit;
version version;
}
join-load-balance;
join-prune-timeout;
mdt {
  data-mdt-reuse;
  group-range multicast-prefix;
  threshold {
    group group-address {
      source source-address {
        rate threshold-rate;
      }
    }
  }
  tunnel-limit limit;
}
}
mvpn {
  autodiscovery {
    inet-mdt;
  }
}
nonstop-routing;
override-interval milliseconds;
propagation-delay milliseconds;
reset-tracking-bit;
rib-group group-name;
rp {
  auto-rp {
    (announce | discovery | mapping);
    (mapping-agent-election | no-mapping-agent-election);
  }
  bidirectional {
    address address {
      group-ranges {
        destination-ip-prefix </prefix-length>;
      }
      hold-time seconds;
      priority number;
    }
  }
  bootstrap {
    family (inet | inet6) {
      export [ policy-names ];
      import [ policy-names ];
      priority number;
    }
  }
  bootstrap-import [ policy-names ];
  bootstrap-export [ policy-names ];
}

```

```

bootstrap-priority number;
dr-register-policy [ policy-names ];
embedded-rp {
    group-ranges {
        destination-ip-prefix </prefix-length>;
    }
    maximum-rps limit;
}
group-rp-mapping {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
}
log-interval seconds;
maximum limit;
threshold value;
}
local {
    family (inet | inet6) {
        address address;
        anycast-pim {
            rp-set {
                address address <forward-msdp-sa>;
            }
            disable;
            local-address address;
        }
        group-ranges {
            destination-ip-prefix </prefix-length>;
        }
        hold-time seconds;
        override;
        priority number;
    }
}
register-limit {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
}
log-interval seconds;
maximum limit;
threshold value;
}
rp-register-policy [ policy-names ];
spt-threshold {
    infinity [ policy-names ];
}
static {
    address address {

```

```

        override;
        version version;
        group-ranges {
            destination-ip-prefix </prefix-length>;
        }
    }
}
rpf-selection {
    group group-address {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
    prefix-list prefix-list-addresses {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
}
sglimit {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
    log-interval seconds;
    maximum limit;
    threshold value;
}
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
tunnel-devices [ mt-fpc/pic/port ];
}

```

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
 [edit protocols],
 [edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced before Junos OS Release 7.4.
family statement introduced in Junos OS Release 9.6.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.

| | |
|---------------------------------|--|
| Description | Enable PIM on the routing device. The remaining statements are explained separately. |
| Default | PIM is disabled on the routing device. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode</i>• <i>Configuring PIM Dense Mode Properties</i>• <i>Configuring PIM Sparse-Dense Mode Properties</i> |

priority (PIM Interfaces)

| | |
|---------------------------------|---|
| Syntax | <code>priority number;</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols <code>pim interface interface-name</code>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim interface interface-name</code>], [edit protocols <code>pim interface interface-name</code>], [edit routing-instances <i>routing-instance-name</i> protocols <code>pim interface interface-name</code>] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Configure the routing device's likelihood to be elected as the designated router. |
| Options | number —Routing device's priority for becoming the designated router. A higher value corresponds to a higher priority. Range: 0 through 4294967295 Default: 1 (Each routing device has an equal probability of becoming the DR.) |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring Interface Priority for PIM Designated Router Selection</i> |

priority (Bootstrap)

| | |
|---------------------------------|--|
| Syntax | <code>priority <i>number</i>;</code> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols <code>pim rp bootstrap</code> (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim rp bootstrap</code> (inet inet6)],</p> <p>[edit protocols <code>pim rp bootstrap</code> (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim rp bootstrap</code> (inet inet6)]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Configure the routing device's likelihood to be elected as the bootstrap router. |
| Options | <p><i>number</i>—Routing device's priority for becoming the bootstrap router. A higher value corresponds to a higher priority.</p> <p>Range: 0 through a 32-bit number</p> <p>Default: 0 (The routing device has the least likelihood of becoming the bootstrap router and sends packets with a priority of 0.)</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring PIM Bootstrap Properties for IPv4</i> • <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i> • bootstrap-priority on page 79 |

priority (PIM RPs)

| | |
|---------------------------------|---|
| Syntax | <code>priority <i>number</i>;</code> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim rp local family (inet inet6)],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6)]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 13.3 for the PTX5000 router.</p> |
| Description | <p>For PIM-SM, configure this routing device's priority for becoming an RP.</p> <p>For bidirectional PIM, configure this RP address' priority for becoming an RP.</p> <p>The bootstrap router uses this field when selecting the list of candidate rendezvous points to send in the bootstrap message. A smaller number increases the likelihood that the routing device or RP address becomes the RP. A priority value of 0 means that bootstrap router can override the group range being advertised by the candidate RP.</p> |
| Options | <p><i>number</i>—Priority for becoming an RP. A lower value corresponds to a higher priority.</p> <p>Range: 0 through 255</p> <p>Default: 1</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring Local PIM RPs</i> in the <i>Multicast Protocols Feature Guide for Routing Devices</i> • <i>Example: Configuring Bidirectional PIM</i> |

promiscuous-mode (Protocols IGMP)

| | |
|---------------------------------|---|
| Syntax | <code>promiscuous-mode;</code> |
| Hierarchy Level | [edit dynamic-profiles <i>profile-name</i> protocols igmp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols igmp <i>interface interface-name</i>], [edit protocols igmp <i>interface interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for dynamic profiles. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify that the interface accepts IGMP reports from hosts on any subnetwork. Note that when enabling promiscuous-mode, all routing devices on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Dynamic IGMP Configuration Overview</i> • <i>Configuring Dynamic DHCP Client Access to a Multicast Network</i> • <i>Accepting IGMP Messages from Remote Subnetworks</i> |

proxy (Multicast VLAN Registration)

| | |
|---------------------------------|---|
| Syntax | <code>proxy source-address <i>ip-address</i>;</code> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 9.6 for EX Series switches. |
| Description | Specify that the VLAN operate in proxy mode. The proxy option is supported only for a VLAN acting as a data-forwarding source. |
| Default | Disabled |
| Options | <code>source-address <i>ip-address</i></code> —IP address of the source VLAN to act as proxy. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Example: Configuring Multicast VLAN Registration on page 33 • Configuring Multicast VLAN Registration (CLI Procedure) on page 60 |

query-interval (Protocols IGMP)

| | |
|---------------------------------|--|
| Syntax | query-interval <i>seconds</i> ; |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp], [edit protocols igmp] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify how often the querier routing device sends general host-query messages. |
| Options | <i>seconds</i> —Time interval. Range: 1 through 1024 Default: 125 seconds |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Modifying the IGMP Host-Query Message Interval</i>• query-last-member-interval (Protocols IGMP) on page 127• query-response-interval (Protocols IGMP) on page 128 |

query-last-member-interval (Protocols IGMP)

| | |
|---------------------------------|--|
| Syntax | query-last-member-interval <i>seconds</i> ; |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp], [edit protocols igmp] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify how often the querier routing device sends group-specific query messages. |
| Options | <i>seconds</i> —Time interval, in fractions of a second or seconds. Range: 0.1 through 0.9, then in 1-second intervals 1 through 999999 Default: 1 second |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Modifying the IGMP Last-Member Query Interval</i> • query-interval (Protocols IGMP) on page 126 • query-response-interval (Protocols IGMP) on page 128 |

query-response-interval (Protocols IGMP)

| | |
|---------------------------------|--|
| Syntax | query-response-interval <i>seconds</i> ; |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp], [edit protocols igmp] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify how long the querier routing device waits to receive a response to a host-query message from a host. |
| Options | seconds —The query response interval must be less than the query interval. Range: 1 through 1024 Default: 10 seconds |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Modifying the IGMP Query Response Interval</i>• query-interval (Protocols IGMP) on page 126• query-last-member-interval (Protocols IGMP) on page 127 |

receiver

| | |
|---------------------------------|---|
| Syntax | <pre> receiver { source-vlans <i>vlan-list</i>; install; } </pre> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>) data-forwarding] |
| Release Information | Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series. |
| Description | <p>Configure a VLAN as a multicast receiver VLAN of the multicast VLAN (MVLAN).</p> <p>The remaining statements are explained separately.</p> |
| Default | Disabled |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Example: Configuring Multicast VLAN Registration on page 33 • Configuring Multicast VLAN Registration (CLI Procedure) on page 60 |

restart-duration (Protocols PIM)

| | |
|---------------------------------|---|
| Syntax | <code>restart-duration seconds;</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim graceful-restart], [edit protocols pim graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols pim graceful-restart] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Configure the duration of the graceful restart interval. |
| Options | seconds —Time that the routing device waits (in seconds) to complete PIM sparse mode graceful restart. Range: 30 through 300 Default: 60 |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring PIM Sparse Mode Graceful Restart</i> |

rib-group (Protocols PIM)

| | |
|---------------------------------|---|
| Syntax | <pre> rib-group { inet <i>group-name</i>; inet6 <i>group-name</i>; } </pre> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Associate a routing table group with PIM. |
| Options | <i>table-name</i> —Name of the routing table. The name must be one that you defined with the rib-groups statement at the [edit routing-options] hierarchy level. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Example: Configuring a Dedicated PIM RPF Routing Table</i> |

robust-count (IGMP Snooping)

| | |
|---------------------------------|--|
| Syntax | <code>robust-count <i>number</i>;</code> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 9.1 for EX Series switches. |
| Description | Configure the number of queries the switch sends before removing a multicast group from the multicast forwarding table. We recommend that the robust count be set to the same value on all multicast routers and switches in the VLAN. |
| Default | The default is the value of the robust-count statement configured for IGMP. The default for the IGMP robust-count statement is 2. |
| Options | <i>number</i> —Number of queries the switch sends before timing out a multicast group. Range: 2 through 10 |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IGMP Snooping (CLI Procedure) on page 39 |

robust-count (Protocols IGMP)

| | |
|---------------------------------|--|
| Syntax | <code>robust-count <i>number</i>;</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp], [edit protocols igmp] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Tune the expected packet loss on a subnet. This factor is used to calculate the group member interval, other querier present interval, and last-member query count. |
| Options | <i>number</i> —Robustness variable. Range: 2 through 10 Default: 2 |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Modifying the IGMP Robustness Variable |

robust-count (MLD Snooping)

| | |
|---------------------------------|--|
| Syntax | <code>robust-count <i>number</i>;</code> |
| Hierarchy Level | <p>[edit protocols mld-snooping vlan (all <i>vlan-name</i>)]</p> <p>[edit routing-instances <i>instance-name</i> protocols mld-snooping vlan <i>vlan-name</i>]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Support at the [edit routing-instances <i>instance-name</i> protocols mld-snooping vlan <i>vlan-name</i>] hierarchy level introduced in Junos OS Release 13.3 for EX Series switches.</p> |
| Description | Configure the number of queries the switch sends before removing a multicast group from the multicast forwarding table. We recommend that the robust count be set to the same value on all multicast routers and switches in the VLAN. |
| Default | The default is the value of the robust-count statement configured for MLD. The default for the MLD robust-count statement is 2. |
| Options | <p>number—Number of queries the switch sends before timing out a multicast group.</p> <p>Range: 2 through 10</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |

rp

```

Syntax  rp {
    auto-rp {
        (announce | discovery | mapping);
        (mapping-agent-election | no-mapping-agent-election);
    }
    bidirectional {
        address address {
            group-ranges {
                destination-ip-prefix </prefix-length>;
            }
            hold-time seconds;
            priority number;
        }
    }
    bootstrap {
        family (inet | inet6) {
            export [ policy-names ];
            import [ policy-names ];
            priority number;
        }
    }
    bootstrap-export [ policy-names ];
    bootstrap-import [ policy-names ];
    bootstrap-priority number;
    dr-register-policy [ policy-names ];
    embedded-rp {
        group-ranges {
            destination-ip-prefix </prefix-length>;
        }
        maximum-rps limit;
    }
    group-rp-mapping {
        family (inet | inet6) {
            log-interval seconds;
            maximum limit;
            threshold value;
        }
    }
    log-interval seconds;
    maximum limit;
    threshold value;
}
local {
    family (inet | inet6) {
        disable;
        address address;
        anycast-pim {
            local-address address;
            address address <forward-msdp-sa>;
            rp-set {
            }
        }
    }
}

```

```

    }
    group-ranges {
        destination-ip-prefix</prefix-length>;
    }
    hold-time seconds;
    override;
    priority number;
}
}
register-limit {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
}
log-interval seconds;
maximum limit;
threshold value;
}
}
register-probe-time register-probe-time;
}
rp-register-policy [ policy-names ];
static {
    address address {
        override;
        version version;
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
    }
}
}
}

```

Hierarchy Level [edit logical-systems *logical-system-name* protocols [pim](#)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
[pim](#)],
[edit protocols [pim](#)],
[edit routing-instances *routing-instance-name* protocols [pim](#)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the routing device as an actual or potential RP. A routing device can be an RP for more than one group.

The remaining statements are explained separately.

Default If you do not include the **rp** statement, the routing device can never become the RP.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding PIM Sparse Mode*

rp-register-policy

Syntax `rp-register-policy [policy-names];`

Hierarchy Level [edit logical-systems *logical-system-name* protocols [pim rp](#)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols [pim rp](#)],
[edit protocols [pim rp](#)],
[edit routing-instances *routing-instance-name* protocols [pim rp](#)]

Release Information Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Apply one or more policies to control incoming PIM register messages.

Options *policy-names*—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Register Message Filters on a PIM RP and DR*
- [dr-register-policy on page 84](#)

rp-set

| | |
|---------------------------------|---|
| Syntax | rp-set { address <i>address</i> <forward-msdp-sa>; } |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim local family (inet inet6) anycast-pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim local family (inet inet6) anycast-pim], [edit protocols pim local family (inet inet6) anycast-pim], [edit routing-instances <i>routing-instance-name</i> protocols pim local family (inet inet6) anycast-pim] |
| Release Information | Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs. The remaining statements are explained separately. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Example: Configuring PIM Anycast With or Without MSDP |

source (Multicast VLAN Registration)

| | |
|---------------------------------|---|
| Syntax | source { groups <i>group-prefix</i> ; } |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>) data-forwarding] |
| Release Information | Statement introduced in Junos OS Release 9.6 for EX Series switches. |
| Description | Configure a VLAN to be a multicast source VLAN (MVLAN). The remaining statement is explained separately. |
| Default | Disabled |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Example: Configuring Multicast VLAN Registration on page 33 • Configuring Multicast VLAN Registration (CLI Procedure) on page 60 |

source (Protocols IGMP)

| | |
|---------------------------------|--|
| Syntax | <code>source <i>ip-address</i> { source-count <i>number</i>; source-increment <i>increment</i>; }</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp interface <i>interface-name</i> static group <i>mcast-group-address</i>], [edit protocols igmp interface <i>interface-name</i> static group <i>mcast-group-address</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify the IP version 4 (IPv4) unicast source address for the multicast group being statically configured on an interface. |
| Options | <i>ip-address</i> —IPv4 unicast address. The remaining statements are explained separately. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Enabling IGMP Static Group Membership</i> |

source-vlans

| | |
|---------------------------------|--|
| Syntax | <code>source-vlans <i>vlan-list</i>;</code> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>) data-forwarding receiver] |
| Release Information | Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series. |
| Description | Specify a list of multicast VLANs (MVLANS) from which this multicast receiver VLAN receives multicast traffic. Either all of these MVLANS must be in proxy mode or none of them can be in proxy mode. |
| Default | Disabled |
| Options | <i>vlan-list</i> —Names of the MVLANS. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring Multicast VLAN Registration on page 33• Configuring Multicast VLAN Registration (CLI Procedure) on page 60 |

spt-threshold

| | |
|---------------------------------|--|
| Syntax | <code>spt-threshold { infinity [<i>policy-names</i>]; }</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim] |
| Release Information | Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | <p>Set the SPT threshold to infinity for a source-group address pair. Last-hop multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or an SPT rooted at the source. By default, last-hop routing devices transition to a direct SPT to the source. You can configure this routing device to set the SPT transition value to infinity to prevent this transition for any source-group address pair.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring the PIM SPT Threshold Policy</i> |

ssm-map (Protocols IGMP)

| | |
|---------------------------------|--|
| Syntax | <code>ssm-map <i>ssm-map-name</i>;</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp interface interface-name], [edit protocols igmp interface interface-name] |
| Release Information | Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Apply an SSM map to an IGMP interface. |
| Options | <i>ssm-map-name</i> —Name of SSM map. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Example: Configuring SSM Mapping |

static (IGMP Snooping)

| | |
|---------------------------------|---|
| Syntax | <code>static { group ip-address; }</code> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>) interface (all <i>interface-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 9.1 for EX Series switches. |
| Description | Statically define multicast groups on an interface. The remaining statement is explained separately. |
| Default | No multicast groups are statically defined. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IGMP Snooping (CLI Procedure) on page 39 • show igmp-snooping membership on page 220 |

static (Protocols PIM)

| | |
|---------------------------------|---|
| Syntax | <pre>static { address address { group-ranges { destination-ip-prefix</prefix-length>; } override; version version; } }</pre> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp], [edit protocols pim rp], [edit routing-instances <i>routing-instance-name</i> protocols pim rp]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Configure static RP addresses. The default static RP address is 224.0.0.0/4. To configure other addresses, include one or more address statements. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.</p> <p>For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Configuring the Static PIM RP Address on the Non-RP Routing Device |

static (Protocols IGMP)

Syntax

```
static {
  group multicast-group-address {
    exclude;
    group-count number;
    group-increment increment;
    source ip-address {
      source-count number;
      source-increment increment;
    }
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp **interface** *interface-name*],
[edit protocols igmp **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Test multicast forwarding on an interface without a receiver host.

The **static** statement simulates IGMP joins on a routing device statically on an interface without any IGMP hosts. It is supported for both IGMPv2 and IGMPv3 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.



NOTE: To prevent joining too many groups accidentally, the **static** statement is not supported with the **interface all** statement.

The remaining statements are explained separately.

Required Privilege routing and trace—To view this statement in the configuration.
Level routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- *Enabling IGMP Static Group Membership*

static (MLD Snooping)

| | |
|---------------------------------|---|
| Syntax | <pre>static { group ip-address; }</pre> |
| Hierarchy Level | [edit protocols mld-snooping vlan (all <i>vlan-name</i>) interface (all <i>interface-name</i>)] [edit routing-instances <i>instance-name</i> protocols mld-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.1 for EX Series switches. Support at the [edit routing-instances <i>instance-name</i> protocols mld-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>] hierarchy level introduced in Junos OS Release 13.3 for EX Series switches. |
| Description | Statically define multicast groups on an interface. The remaining statement is explained separately. |
| Default | No multicast groups are statically defined. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |

traceoptions (Protocols PIM)

| | |
|----------------------------|--|
| Syntax | <pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre> |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols pim],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],</p> <p>[edit protocols pim],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Configure PIM tracing options.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p> |
| Default | The default PIM trace options are those inherited from the routing protocol's traceoptions statement included at the [edit routing-options] hierarchy level. |
| Options | <p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place tracing output in the pim-log file.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <p>PIM Tracing Flags</p> <ul style="list-style-type: none"> assert—Assert messages bidirectional-df-election—Bidirectional PIM designated-forwarder (DF) election events |

- **bootstrap**—Bootstrap messages
- **cache**—Packets in the PIM sparse mode routing cache
- **graft**—Graft and graft acknowledgment messages
- **hello**—Hello packets
- **join**—Join messages
- **mt**—Multicast tunnel messages
- **nsr-synchronization**—Nonstop active routing (NSR) synchronization messages
- **packets**—All PIM packets
- **prune**—Prune messages
- **register**—Register and register stop messages
- **rp**—Candidate RP advertisements
- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 0 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

| | |
|------------------------------|---|
| Required Privilege | routing and trace—To view this statement in the configuration. |
| Level | routing-control and trace-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring PIM Trace Options</i> • <i>Tracing DVMRP Protocol Traffic</i> • <i>Tracing MSDP Protocol Traffic</i> • <i>Configuring PIM Trace Options</i> |

traceoptions (Protocols IGMP)

| | |
|----------------------------|---|
| Syntax | <pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp], [edit protocols igmp] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | <p>Configure IGMP tracing options.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p> <p>To trace the paths of multicast packets, use the mtrace command.</p> |
| Default | The default IGMP trace options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level. |
| Options | <p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place tracing output in the file igmp-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <p>IGMP Tracing Flags</p> <ul style="list-style-type: none">• leave—Leave group messages (for IGMP version 2 only).• mtrace—Mtrace packets. Use the mtrace command to troubleshoot the software. |

- **packets**—All IGMP packets.
- **query**—IGMP membership query messages, including general and group-specific queries.
- **report**—Membership report messages.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

| | |
|------------------------------|--|
| Required Privilege | routing and trace—To view this statement in the configuration. |
| Level | routing-control and trace-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Tracing IGMP Protocol Traffic</i> |

traceoptions (IGMP Snooping)

| | |
|----------------------------|--|
| Syntax | <pre> traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <replace> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier>; } </pre> |
| Hierarchy Level | [edit protocols igmp-snooping] |
| Release Information | Statement introduced in Junos OS Release 9.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. |
| Description | Define tracing operations for IGMP snooping. |
| Default | The traceoptions feature is disabled by default. |
| Options | <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files, including the active trace file. When a trace file reaches its maximum size, its contents are archived into a compressed file named <i>filename.0</i> and the trace file is emptied. When the trace file reaches its maximum size again, the <i>filename.0</i> archive file is renamed <i>filename.1</i> and a new <i>filename.0</i> archive file is created from the contents of the trace file. This process continues until the maximum number of trace files is reached, at which point the system starts overwriting the oldest archive file each time the trace file is archived. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 10 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> • all—All tracing operations. • general—Trace general IGMP snooping protocol events. • krt—Trace communication over routing socket. • leave—Trace leave group messages (IGMPv2 and IGMPv3 only). • nexthop—Trace nexthop-related events. • normal—Trace normal IGMP snooping protocol events. If you do not specify this flag, only unusual or abnormal operations are traced. • packets—Trace all IGMP packets. • policy—Trace policy processing. • query—Trace IGMP membership query messages. |

- **report**—Trace membership report messages.
- **route**—Trace routing information.
- **state**—Trace IGMP state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.
- **vlan**—Trace VLAN-related events.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers per flag:

- **detail**—Provide detailed trace information
- **disable**—Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

no-stamp—(Optional) Omit the timestamp at the beginning of each line in the trace file.

no-world-readable—(Optional) Restrict file access to the user who created the file.

replace—(Optional) Replace an existing trace file if there is one. If you do not include this option, tracing output is appended to an existing trace file.

size size —(Optional) Maximum size of each trace file, in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is zipped and renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum size, you also must specify a maximum number of files with the **files** option.

Syntax: *x* to specify bytes, *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10240 through 4294967295 bytes

Default: 128 KB

world-readable—(Optional) Allow unrestricted file access.

| | |
|---------------------------------|---|
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
|---------------------------------|---|

| | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none">• Configuring IGMP Snooping Tracing Operations (CLI Procedure) on page 48 |
|------------------------------|---|

traceoptions (MLD Snooping)

| | |
|----------------------------|--|
| Syntax | <pre> traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <replace> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>>; } </pre> |
| Hierarchy Level | <p>[edit protocols mld-snooping]</p> <p>[edit protocols mld-snooping <i>vlan</i> <i>vlan-name</i>]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Support at the [edit protocols mld-snooping <i>vlan</i> <i>vlan-name</i>] hierarchy level introduced in Junos OS Release 13.3 for EX Series switches.</p> |
| Description | Define tracing operations for MLD snooping. |
| Default | The traceoptions feature is disabled by default. |
| Options | <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <i>/var/log</i>.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files, including the active trace file. When a trace file reaches its maximum size, its contents are archived into a compressed file named <i>filename.0</i> and the trace file is emptied. When the trace file reaches its maximum size again, the <i>filename.0</i> archive file is renamed <i>filename.1</i> and a new <i>filename.0</i> archive file is created from the contents of the trace file. This process continues until the maximum number of trace files is reached, at which point the system starts overwriting the oldest archive file each time the trace file is archived. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 10 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> • all—All tracing operations. • client-notification—(EX9200 switches only) Trace client notifications. • general—Trace general MLD snooping protocol events. • group—(EX9200 switches only) Trace group operations. • host-notification—(EX9200 switches only) Trace host notifications. • krt—(All EX Series switches except EX9200) Trace communication over routing socket. • leave—Trace leave group messages. • nexthop—(All EX Series switches except EX9200) Trace events related to next-hops. |

- **normal**—Trace normal MLD snooping protocol events. If you do not specify this flag, only unusual or abnormal operations are traced.
- **packets**—Trace all MLD packets.
- **policy**—Trace policy processing.
- **query**—Trace MLD membership query messages.
- **report**—Trace membership report messages.
- **route**—Trace routing information.
- **state**—Trace MLD state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.
- **vlan**—(All EX Series switches except EX9200) Trace VLAN-related events.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers per flag:

- **detail**—Provide detailed trace information
- **disable**—Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

no-stamp—(All EX Series switches except EX9200) (Optional) Omit the timestamp at the beginning of each line in the trace file.

no-world-readable—(Optional) Restrict file access to the user who created the file.

replace—(All EX Series switches except EX9200) (Optional) Replace an existing trace file if there is one. If you do not include this option, tracing output is appended to an existing trace file.

size size —(Optional) Maximum size of each trace file, in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is zipped and renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum size, you also must specify a maximum number of files with the **files** option.

Syntax: *x* to specify bytes, *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10240 through 4294967295 bytes

Default: 128 KB

world-readable—(Optional) Allow unrestricted file access.

| | |
|---------------------------|---|
| Required Privilege | routing—To view this statement in the configuration. |
| Level | routing-control—To add this statement to the configuration. |

Related
Documentation

version (Protocols IGMP)

| | |
|--------------------------|--|
| Syntax | <code>version version;</code> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> protocols igmp interface interface-name], [edit protocols igmp interface interface-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify the version of IGMP. |
| Options | version —IGMP version number. Range: 1, 2, or 3 Default: IGMP version 2 |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Changing the IGMP Version</i> |

version (IGMP Snooping)

| | |
|---------------------------------|--|
| Syntax | <code>version number;</code> |
| Hierarchy Level | [edit protocols igmp-snooping vlan (all <i>vlan-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Specify the IGMP version for the IGMP general query that the switch sends to hosts when an interface comes up. The configured IGMP version affects only the version of the general queries sent by a switch. It does not affect the version of IGMP messages that the switch can snoop. For example, If the switch is configured for IGMP version 1 (IGMPv1), it can snoop IGMPv2 and IGMPv3 messages. |
| Default | If you do not configure the version statement, the default is IGMPv2. |
| Options | version —IGMP version number. Range: 1 and 2. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IGMP Snooping (CLI Procedure) on page 39• Configuring IGMP Snooping |

version (MLD Snooping)

| | |
|---------------------------------|---|
| Syntax | <code>version version;</code> |
| Hierarchy Level | [edit protocols mld-snooping vlan (all <i>vlan-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 12.1 for EX Series switches. |
| Description | Specify the MLD version for the MLD general query the switch sends to hosts when an interface comes up. The configured MLD version affects only the version of the general queries sent by a switch: it does not affect the version of MLD messages that the switch can snoop. If the switch is configured for MLD version 1 (MLDv1), it can snoop MLDv2 messages and vice versa. |
| Default | MLDv1 is the default. |
| Options | version —MLD version number. Values: 1 or 2 |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |

version (PIM)

| | |
|---------------------------------|---|
| Syntax | <code>version version;</code> |
| Hierarchy Level | <code>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols pim rp static address <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> pim interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> pim rp static address <i>address</i>],</code> <code>[edit protocols pim interface <i>interface-name</i>],</code> <code>[edit protocols pim rp static address <i>address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim rp static address <i>address</i>]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify the version of PIM. |
| Options | version —PIM version number. Range: 1 or 2 Default: PIMv1 for rendezvous point (RP) mode (at the <code>[edit protocols pim rp static address <i>address</i>]</code> hierarchy level). PIMv2 for interface mode (at the <code>[edit protocols pim interface <i>interface-name</i>]</code> hierarchy level). |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Enabling PIM Sparse Mode</i>• <i>Configuring PIM Dense Mode Properties</i>• <i>Configuring PIM Sparse-Dense Mode Properties</i> |

vlan (IGMP Snooping)

```
Syntax  vlan (all | vlan-name) {
        data-forwarding {
            source {
                groups group-prefix;
            }
            receiver {
                source-vlans vlan-list;
                install;
            }
        }
        disable;
        immediate-leave;
        interface (all | interface-name) {
            multicast-router-interface;
            static {
                group ip-address;
            }
        }
        proxy {
            source-address ip-address;
        }
        robust-count number;
        version version;
    }
```

Hierarchy Level [edit protocols **igmp-snooping**]

Release Information Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement updated with enhanced ? (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.

Description Configure IGMP snooping parameters for a VLAN.

When the **vlan** configuration statement is used without the **disable** statement, IGMP snooping is enabled on the specified VLAN or on all VLANs.



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

Default If the **vlan** statement is not included in the configuration, IGMP snooping is disabled.

- Options**
- **all**—All VLANs on the switch
 - ***vlan-name***—Name of a VLAN.



TIP: When you configure IGMP snooping parameters using the **vlan all** statement, any VLAN that is not individually configured for IGMP snooping

inherits the `vlan all` configuration. Any VLAN that is individually configured for IGMP snooping, on the other hand, inherits none of its configuration from `vlan all`. Any parameters that are not explicitly defined for the individual VLAN assume their default values, not the values specified in the `vlan all` configuration.

For example, in the following configuration:

```
protocols {
  igmp-snooping {
    vlan all {
      robust-count 8;
    }
    vlan employee {
      interface ge-0/0/8.0 {
        static {
          group 239.0.10.3
        }
      }
    }
  }
}
```

all VLANs, except `employee`, have a robust count of 8. Because `employee` has been individually configured, its robust count value is not determined by the value set under `vlan all`. Instead, its robust count is the default value of 2.

The remaining statements are explained separately.

| | |
|---------------------------|---|
| Required Privilege | routing—To view this statement in the configuration. |
| Level | routing-control—To add this statement to the configuration. |

| | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring IGMP Snooping on EX Series Switches on page 27• Configuring IGMP Snooping (CLI Procedure) on page 39• show igmp-snooping vlans on page 227 |
|------------------------------|---|

vlan (MLD Snooping)

| | |
|----------------------------|---|
| Syntax | <pre> vlan (all <i>vlan-name</i>) { disable; immediate-leave; interface (all <i>interface-name</i>) { group-limit <i>limit</i>; host-only-interface; immediate-leave; multicast-router-interface; static { group <i>ip-address</i> { source <i>ip-address</i>; } } } qualified-vlan; query-interval <i>seconds</i>; query-last-member-interval <i>seconds</i>; query-response-interval <i>seconds</i>; robust-count <i>number</i>; traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier>; } version <i>version</i>; } </pre> |
| Hierarchy Level | <p>[edit protocols mld-snooping]</p> <p>[edit routing-instances <i>instance-name</i> protocols mld-snooping]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Support at the [edit routing-instances <i>instance-name</i> protocols mld-snooping] hierarchy introduced in Junos OS Release 13.3 for EX Series switches.</p> <p>Support for the qualified-vlan, query-interval, query-last-member-interval, query-response-interval, and traceoptions statements introduced in Junos OS Release 13.3 for EX Series switches.</p> |
| Description | <p>Configure MLD snooping parameters for a VLAN.</p> <p>When the vlan configuration statement is used without the disable statement, MLD snooping is enabled on the specified VLAN or on all VLANs.</p> |
| Default | <p>If the vlan statement is not included in the configuration, MLD snooping is disabled.</p> |
| Options | <p>all—(All EX Series switches except EX9200) Configure MLD snooping parameters for all VLANs on the switch.</p> <p><i>vlan-name</i>—Configure MLD snooping parameters for the specified VLAN.</p> |



TIP: When you configure MLD snooping parameters using the `vlan all` statement, any VLAN that is not individually configured for MLD snooping inherits the `vlan all` configuration. Any VLAN that is individually configured for MLD snooping, on the other hand, inherits none of its configuration from `vlan all`. Any parameters that are not explicitly defined for the individual VLAN assume their default values, not the values specified in the `vlan all` configuration.

For example, in the following configuration:

```
protocols {
  mld-snooping {
    vlan all {
      robust-count 8;
    }
    vlan employee {
      interface ge-0/0/8.0 {
        static {
          group ff1e::1;
        }
      }
    }
  }
}
```

all VLANs, except `employee`, have a robust count of 8. Because `employee` has been individually configured, its robust count value is not determined by the value set under `vlan all`. Instead, its robust count is the default value of 2.

The remaining statements are explained separately.

| | |
|---------------------------|---|
| Required Privilege | routing—To view this statement in the configuration. |
| Level | routing-control—To add this statement to the configuration. |

| | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none"> • Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |
|------------------------------|---|

PART 3

Administration

- [Routine Monitoring on page 165](#)
- [Operational Commands on page 175](#)

CHAPTER 7

Routine Monitoring

- [Monitoring IGMP Snooping on page 165](#)
- [Verifying IGMP Snooping \(CLI Procedure\) on page 168](#)
- [Verifying MLD Snooping \(CLI Procedure\) on page 170](#)

Monitoring IGMP Snooping

Purpose



NOTE: This topic applies only to the J-Web Application package.

Use the monitoring feature to view status and information about IGMP snooping configuration on your EX Series switch.

Action

To display IGMP snooping details in the J-Web interface, select **Monitor > Switching > IGMP Snooping**.

To display IGMP snooping details in the CLI, enter the following commands:

- `show igmp-snooping route`
- `show igmp-snooping statistics`
- `show igmp-snooping vlans`



NOTE: On EX4300 switches, to display IGMP snooping details in the CLI, enter the following commands:

- `show igmp snooping interface`
- `show igmp snooping statistics`
- `show multicast snooping next-hops`
- `show multicast snooping route`

Meaning

[Table 6 on page 166](#) summarizes the IGMP snooping details displayed.

Table 6: Summary of IGMP Snooping Output Fields

| Field | Values |
|---|--|
| IGMP Snooping Monitor | |
| VLAN | The VLAN for which IGMP snooping is enabled. |
| Interfaces | Indicates the interfaces configured as switching interfaces that are associated with the multicast router. |
| Groups | Indicates the number of the multicast groups learned by the VLAN. |
| Learning Domain | Learning domain for snooping. |
| NOTE: This option is supported only on EX4300 switches. | |
| Query Interval | Frequency (in seconds) with which the router sends membership queries when it is the querier. |
| NOTE: This option is supported only on EX4300 switches. | |
| Query Response Interval | Time (in seconds) that the router waits for a response to a general query. |
| NOTE: This option is supported only on EX4300 switches. | |
| Last Member Query Interval | Time (in seconds) that the router waits for a report in response to a group-specific query. |
| NOTE: This option is supported only on EX4300 switches. | |
| Robustness Count | Number of times the router retries a query. |
| NOTE: This option is supported only on EX4300 switches. | |
| MRouters | Specifies the multicast router. |
| NOTE: This option is not supported on EX4300 switches. | |
| Receivers | Specifies the multicast receiver. |
| NOTE: This option is not supported on EX4300 switches. | |
| Interface Details | |
| Interfaces | Name of the interface. |
| NOTE: This option is only supported on EX4300 switches. | |
| State | Operating state of the interface. Values are Up or Down. |
| NOTE: This option is only supported on EX4300 switches. | |
| Group count | Number of groups on the interface. |
| NOTE: This option is only supported on EX4300 switches. | |

Table 6: Summary of IGMP Snooping Output Fields (*continued*)

| Field | Values |
|--|--|
| Immediate Leave NOTE: This option is only supported on EX4300 switches. | State of immediate leave. Values are On or Off. |
| Router Interface NOTE: This option is only supported on EX4300 switches. | Router interfaces that are part of this learning domain. |
| IGMP Route Information | |
| VLAN NOTE: This option is not supported on EX4300 switches. | The VLAN for which IGMP snooping is enabled. |
| Group NOTE: This option is not supported on EX4300 switches. | Indicates the multicast groups learned by the VLAN. |
| Next-Hop NOTE: This option is not supported on EX4300 switches. | Specifies the next hop assigned by the switch after performing the route lookup. |
| Statistics | |
| Packets per Vlan NOTE: This option is supported only on EX4300 switches. | Displays the number of packets sent or received, or the number of received errors. |
| Global NOTE: This option is supported only on EX4300 switches. | Displays the statistics name and statistics count. The statistics names are: <ul style="list-style-type: none"> • Bad Length • Bad Checksum • Bad Receive If • Rx non-local • Timed out |
| Multicast Snooping Nexthops | |
| ID NOTE: This option is supported only on EX4300 switches. | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine. |
| Reference Count NOTE: This option is supported only on EX4300 switches. | Number of cache entries that are using this next hop. |
| Kernel Reference Count NOTE: This option is supported only on EX4300 switches. | Kernel reference count for the next hop. |

Table 6: Summary of IGMP Snooping Output Fields (*continued*)

| Field | Values |
|----------------------|---|
| Downstream Interface | Interface names associated with each multicast next-hop ID. |

NOTE: This option is supported only on EX4300 switches.

Related Documentation

- [show igmp-snooping route on page 223](#)
- [show igmp-snooping statistics on page 225](#)
- [show igmp-snooping vlans on page 227](#)
- [Configuring IGMP Snooping \(CLI Procedure\) on page 39](#)
- [Example: Configuring IGMP Snooping on EX Series Switches on page 27](#)

Verifying IGMP Snooping (CLI Procedure)

Internet Group Management Protocol (IGMP) snooping constrains the flooding of IPv4 multicast traffic on VLANs on a switch. This topic describes how to verify IGMP snooping operation on the switch.

It covers:

- [Verifying IGMP Snooping Memberships on page 168](#)
- [Viewing IGMP Snooping Statistics on page 169](#)
- [Viewing IGMP Snooping Routing Information on page 169](#)

Verifying IGMP Snooping Memberships

Purpose Determine group memberships, multicast-router interfaces, host IGMP versions, and the current values of timeout counters.

Action Enter the following command:

```
user@switch> show igmp-snooping membership detail
VLAN: vlan2 Tag: 2 (Index: 3)
  Router interfaces:
    ge-1/0/0.0 dynamic Uptime: 00:14:24 timeout: 253
  Group: 225.0.0.1
    ge-1/0/17.0 259 Last reporter: 13.0.0.90 Receiver count: 1
    Uptime: 00:00:19 timeout: 259 Flags: <V3-hosts>
    Include source: 10.2.11.5, 10.2.11.12
```

Meaning The switch has multicast membership information for one VLAN on the switch, **vlan2**. IGMP snooping might be enabled on other VLANs, but the switch does not have any multicast membership information for them. The following information is provided:

- Information on the multicast-router interfaces for the VLAN—in this case, **ge-1/0/0.0**. The multicast-router interface has been learned by IGMP snooping, as indicated by the dynamic value. The timeout value shows how many seconds from now the interface

will be removed from the multicast forwarding table if the switch does not receive IGMP queries or Protocol Independent Multicast (PIM) updates on the interface.

- Information about the group memberships for the VLAN:
 - Currently, the VLAN has membership in only one multicast group, **225.0.0.1**.
 - The host or hosts that have reported membership in the group are on interface **ge-1/0/17.0**. The last host that reported membership in the group has address **13.0.0.90**. The number of hosts belonging to the group on the interface is shown in the Receiver count field, which is displayed only when host tracking is enabled if immediate leave is configured on the VLAN.
 - The Uptime field shows that the multicast group has been active on the interface for 19 seconds. The interface group membership will time out in 259 seconds if no hosts respond to membership queries during this interval. The Flags field shows the lowest version of IGMP used by a host that is currently a member of the group, which in this case is IGMP version 3 (IGMPv3).
 - Because the interface has IGMPv3 hosts on it, the source addresses from which the IGMPv3 hosts want to receive group multicast traffic are shown (addresses **10.2.11.5** and **10.2.11.12**). The timeout value for the interface group membership is derived from the largest timeout value for all sources addresses for the group.

Viewing IGMP Snooping Statistics

Purpose Display IGMP snooping statistics, such as number of IGMP queries, reports, and leaves received and how many of these IGMP messages contained errors.

Action Enter the following command:

```
user@switch> show igmp-snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 0
```

| IGMP Type | Received | Transmitted | Recv Errors |
|-----------|----------|-------------|-------------|
| Queries: | 74295 | 0 | 0 |
| Reports: | 18148423 | 0 | 16333523 |
| Leaves: | 0 | 0 | 0 |
| Other: | 0 | 0 | 0 |

Meaning The output shows how many IGMP messages of each type—**Queries**, **Reports**, **Leaves**—the switch received or transmitted on interfaces on which IGMP snooping is enabled. For each message type, it also shows the number of IGMP packets the switch received that had errors—for example, packets that do not conform to the IGMPv1, IGMPv2, or IGMPv3 standards. If the **Recv Errors** count increases, verify that the hosts are compliant with IGMP standards. If the switch is unable to recognize the IGMP message type for a packet, it counts the packet under **Receive unknown**.

Viewing IGMP Snooping Routing Information

Purpose Display the next-hop information maintained in the multicast forwarding table.

Action Enter the following command:

```
user@switch> show igmp-snooping route detail
VLAN          Group          Next-hop
v100          224.0.0.0, *    1323
               Interfaces: ge-0/0/0.0
VLAN          Group          Next-hop
v100          226.0.0.1, *    1322
               Interfaces: ge-0/0/0.0, ge-0/0/1.0, ge-0/0/47.0
```

Meaning The output shows the next-hop interfaces for a given multicast group on a VLAN. For example, route **226.0.0.1** on **v100** has next-hop interfaces **ge-0/0/0.0**, **ge-0/0/1.0**, and **ge-0/0/47.0**.

- Related Documentation**
- [clear igmp-snooping membership on page 182](#)
 - [clear igmp-snooping statistics on page 183](#)
 - [Example: Configuring IGMP Snooping on EX Series Switches on page 27](#)
 - [Configuring IGMP Snooping \(CLI Procedure\) on page 39](#)

Verifying MLD Snooping (CLI Procedure)

Multicast Listener Discovery (MLD) snooping constrains the flooding of IPv6 multicast traffic on VLANs on a switch. This topic describes how to verify MLD snooping operation on the switch.

It covers:

- [Verifying MLD Snooping Memberships on page 170](#)
- [Verifying MLD Snooping VLANs on page 171](#)
- [Viewing MLD Snooping Statistics on page 172](#)
- [Viewing MLD Snooping Routing Information on page 172](#)

Verifying MLD Snooping Memberships

Purpose Determine group memberships, multicast-router interfaces, host MLD versions, and the current values of timeout counters.

Action Enter the following command:

```
user@switch> show mld-snooping membership detail
VLAN: mld-vlan Tag: 100 (Index: 3)
  Router interfaces:
    ge-1/0/0.0 dynamic Uptime: 00:14:24 timeout: 253
  Group: ff1e::2010
    ge-1/0/30.0 Timeout: 180 Flags: <V2-hosts>
    Last reporter: fe80::2020:1:1:3
    Include source: 2020:1:1:1::2
    Include source: 2020:1:1:1::5
```


Meaning The switch has multicast membership information for one VLAN on the switch, **mld-vlan**. MLD snooping might be enabled on other VLANs, but the switch does not have any multicast membership information for them. The following information is provided:

- Information on the multicast-router interfaces for the VLAN—in this case, **ge-1/0/0.0**. The multicast-router interface has been learned by MLD snooping, as indicated by **dynamic**. The **timeout** value shows how many seconds from now the interface will be removed from the multicast forwarding table if the switch does not receive MLD queries or Protocol Independent Multicast (PIM) updates on the interface.
- Information about the group memberships for the VLAN:
 - Currently, the VLAN has membership in only one multicast group, **ff1e::2010**.
 - The host or hosts that have reported membership in the group are on interface **ge-1/0/30.0**. The interface group membership will time out in 180 seconds if no hosts respond to membership queries during this interval. The flags field shows the lowest version of MLD used by a host that is currently a member of the group, which in this case is MLD version 2 (MLDv2).
 - The last host that reported membership in the group has address **fe80::2020:1:1:3**.
 - Because interface has MLDv2 hosts on it, the source addresses from which the MLDv2 hosts want to receive group multicast traffic are shown (addresses **2020:1:1:1::2** and **2020:1:1:1::5**). The **timeout** value for the interface group membership is derived from the largest timeout value for all sources addresses for the group.

Verifying MLD Snooping VLANs

Purpose Verify that MLD snooping is enabled on a VLAN and display MLD snooping information for each VLAN on which MLD snooping is enabled.

Action Enter the following command:

```
user@switch> show mld-snooping vlans detail
VLAN: v10, Tag: 10
  Interface: ge-1/0/0.0, tagged, Groups: 0, Router
  Interface: ge-1/0/30.0, untagged, Groups: 1
  Interface: ge-12/0/30.0, untagged, Groups: 0
VLAN: v20, Tag: 20
  Interface: ge-1/0/0.0, tagged, Groups: 0, Router
  Interface: ge-1/0/31.0, untagged, Groups: 0
  Interface: ge-12/0/31.0, untagged, Groups: 1
```

Meaning MLD snooping is configured on two VLANs on the switch: **v10** and **v20**. Each interface in each VLAN is listed and the following information is provided:

- Whether the interface is a trunk (**tagged**) or access (**untagged**) interface.
- How many multicast groups the interface belongs to.
- Whether the interface is a multicast-router interface (**Router**).

Viewing MLD Snooping Statistics

Purpose Display MLD snooping statistics, such as number of MLD queries, reports, and leaves received and how many of these MLD messages contained errors.

Action Enter the following command:

```
user@switch> show mld-snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 0
```

| MLD Type | Received | Transmitted | Recv Errors |
|----------|----------|-------------|-------------|
| Queries: | 74295 | 0 | 0 |
| Reports: | 18148423 | 0 | 16333523 |
| Leaves: | 0 | 0 | 0 |
| Other: | 0 | 0 | 0 |

Meaning The output shows how many MLD messages of each type—**Queries**, **Reports**, **Leaves**—the switch received or transmitted on interfaces on which MLD snooping is enabled. For each message type, it also shows the number of MLD packets the switch received that had errors—for example, packets that do not conform to the MLDv1 or MLDv2 standards. If the **Recv Errors** count increases, verify that the hosts are compliant with MLDv1 or MLDv2 standards. If the switch is unable to recognize the MLD message type for a packet, it counts the packet under **Receive unknown**.

Viewing MLD Snooping Routing Information

Purpose Display the next-hop information maintained in the multicast forwarding table.

Action Enter the following command:

```
user@switch> show mld-snooping route detail
VLAN          Group          Next-hop
mld-vlan      ::0000:2010    1323
Interfaces: ge-1/0/30.0, ge-1/0/33.0
VLAN          Group          Next-hop
mld-vlan      ff00::         1317
Interfaces: ge-1/0/0.0, ge-1/0/33.0
VLAN          Group          Next-hop
mld-vlan      ::0000:0000    1317
Interfaces: ge-1/0/0.0
VLAN          Group          Next-hop
mld-vlan1     ::0000:2010    1324
Interfaces: ge-12/0/31.0
VLAN          Group          Next-hop
mld-vlan1     ff00::         1318
Interfaces: ae200.0
VLAN          Group          Next-hop
mld-vlan1     ::0000:0000    1318
Interfaces: ae200.0
```

Meaning The output shows the next-hop interfaces for a given multicast group on a VLAN. Only the last 32 bits of the group address are shown because the switch uses only these bits in determining multicast routes. For example, route **::0000:2010** on **mld-vlan** has next-hop interfaces **ge-1/0/30.0** and **ge-1/0/33.0**.

- Related Documentation**
- [clear mld-snooping membership on page 184](#)
 - [clear mld-snooping statistics on page 185](#)
 - [Example: Configuring MLD Snooping on page 30](#)
 - [Configuring MLD Snooping on a VLAN \(CLI Procedure\) on page 50](#)

CHAPTER 8

Operational Commands

- clear igmp membership
- clear igmp statistics
- clear igmp-snooping membership
- clear igmp-snooping statistics
- clear mld-snooping membership
- clear mld-snooping statistics
- clear multicast bandwidth-admission
- clear multicast scope
- clear multicast sessions
- clear multicast statistics
- clear pim join
- clear pim register
- clear pim statistics
- mtrace
- mtrace from-source
- mtrace monitor
- mtrace to-gateway
- show igmp group
- show igmp interface
- show igmp statistics
- show igmp-snooping membership
- show igmp-snooping route
- show igmp-snooping statistics
- show igmp-snooping vlans
- show mld-snooping membership
- show mld-snooping route
- show mld-snooping statistics
- show mld-snooping vlans

- `show multicast flow-map`
- `show multicast interface`
- `show multicast mrinfo`
- `show multicast next-hops`
- `show multicast pim-to-igmp-proxy`
- `show multicast pim-to-mld-proxy`
- `show multicast route`
- `show multicast rpf`
- `show multicast scope`
- `show multicast sessions`
- `show multicast usage`
- `show pim bootstrap`
- `show pim interfaces`
- `show pim join`
- `show pim neighbors`
- `show pim rps`
- `show pim source`
- `show pim statistics`

clear igmp membership

| | |
|---|--|
| List of Syntax | Syntax on page 177 Syntax (EX Series Switch and the QFX Series) on page 177 |
| Syntax | <pre>clear igmp membership <group address-range> <interface interface-name> <logical-system (all logical-system-name)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>clear igmp membership <group address-range> <interface interface-name></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Clear Internet Group Management Protocol (IGMP) group members. |
| Options | <p>none—Clear all IGMP members on all interfaces and for all address ranges.</p> <p>group address-range—(Optional) Clear all IGMP members that are in a particular address range. An example of a range is 224.2/16. If you omit the destination prefix length, the default is /32.</p> <p>interface interface-name—(Optional) Clear all IGMP group members on an interface.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show igmp group on page 209 • show igmp interface on page 213 |
| List of Sample Output | clear igmp membership on page 177 clear igmp membership interface on page 178 clear igmp membership group on page 179 |
| Output Fields | See show igmp group for an explanation of output fields. |

Sample Output

clear igmp membership

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

```

user@host> show igmp group
Interface      Group           Last Reported   Timeout
so-0/0/0       224.2.127.253  10.1.128.1      186
so-0/0/0       224.2.127.254  10.1.128.1      186
so-0/0/0       239.255.255.255 10.1.128.1      187
so-0/0/0       224.1.127.255   10.1.128.1      188
local          224.0.0.6       (null)           0
local          224.0.0.5       (null)           0
local          224.2.127.254   (null)           0
local          239.255.255.255 (null)           0
local          224.0.0.2       (null)           0
local          224.0.0.13      (null)           0

```

```

user@host> clear igmp membership
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0

```

```

user@host> show igmp group
Interface      Group           Last Reported   Timeout
local          224.0.0.6       (null)           0
local          224.0.0.5       (null)           0
local          224.2.127.254   (null)           0
local          239.255.255.255 (null)           0
local          224.0.0.2       (null)           0
local          224.0.0.13      (null)           0

```

clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```

user@host> show igmp group
Interface      Group           Last Reported   Timeout
so-0/0/0       224.2.127.253  10.1.128.1      210
so-0/0/0       239.255.255.255 10.1.128.1      210
so-0/0/0       224.1.127.255   10.1.128.1      215
so-0/0/0       224.2.127.254   10.1.128.1      216
local          224.0.0.6       (null)           0
local          224.0.0.5       (null)           0
local          224.2.127.254   (null)           0
local          239.255.255.255 (null)           0
local          224.0.0.2       (null)           0
local          224.0.0.13      (null)           0

```

```

user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0

```

```

user@host> show igmp group
Interface      Group           Last Reported   Timeout
local          224.0.0.6       (null)           0
local          224.0.0.5       (null)           0
local          224.2.127.254   (null)           0
local          239.255.255.255 (null)           0
local          224.0.0.2       (null)           0
local          224.0.0.13      (null)           0

```


clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
```

| Interface | Group | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0 | 224.2.127.253 | 10.1.128.1 | 210 |
| so-0/0/0 | 239.255.255.255 | 10.1.128.1 | 210 |
| so-0/0/0 | 224.1.127.255 | 10.1.128.1 | 215 |
| so-0/0/0 | 224.2.127.254 | 10.1.128.1 | 216 |
| local | 224.0.0.6 | (null) | 0 |
| local | 224.0.0.5 | (null) | 0 |
| local | 224.2.127.254 | (null) | 0 |
| local | 239.255.255.255 | (null) | 0 |
| local | 224.0.0.2 | (null) | 0 |
| local | 224.0.0.13 | (null) | 0 |

```
user@host> clear igmp membership group 239.225/16
```

```
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-1/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-2/0/0
```

```
user@host> show igmp group
```

| Interface | Group | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0 | 224.1.127.255 | 10.1.128.1 | 231 |
| so-0/0/0 | 224.2.127.254 | 10.1.128.1 | 233 |
| so-0/0/0 | 224.2.127.253 | 10.1.128.1 | 236 |
| local | 224.0.0.6 | (null) | 0 |
| local | 224.0.0.5 | (null) | 0 |
| local | 224.2.127.254 | (null) | 0 |
| local | 239.255.255.255 | (null) | 0 |
| local | 224.0.0.2 | (null) | 0 |
| local | 224.0.0.13 | (null) | 0 |

clear igmp statistics

| | |
|-----------------------------|--|
| List of Syntax | Syntax on page 180 Syntax (EX Series Switches) on page 180 |
| Syntax | <code>clear igmp statistics</code> <code><interface <i>interface-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Syntax (EX Series Switches) | <code>clear igmp statistics</code> <code><interface <i>interface-name</i>></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Clear Internet Group Management Protocol (IGMP) statistics. |
| Options | none —Clear IGMP statistics on all interfaces. interface <i>interface-name</i> —(Optional) Clear IGMP statistics for the specified interface only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show igmp statistics on page 217 |
| List of Sample Output | clear igmp statistics on page 180 |
| Output Fields | See show igmp statistics for an explanation of output fields. |

Sample Output

clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        8883         459      0
V1 Membership Report    0            0        0
DVMRP                   19784        35476    0
PIM V1                  18310        0        0
Cisco Trace             0            0        0
V2 Membership Report    0            0        0
Group Leave             0            0        0
```

| | | | |
|-------------------------------------|---|---|---|
| Mtrace Response | 0 | 0 | 0 |
| Mtrace Request | 0 | 0 | 0 |
| Domain Wide Report | 0 | 0 | 0 |
| V3 Membership Report | 0 | 0 | 0 |
| Other Unknown types | | | 0 |
| IGMP v3 unsupported type | | | 0 |
| IGMP v3 source required for SSM | | | 0 |
| IGMP v3 mode not applicable for SSM | | | 0 |

| | |
|------------------------|------|
| IGMP Global Statistics | |
| Bad Length | 0 |
| Bad Checksum | 0 |
| Bad Receive If | 0 |
| Rx non-local | 1227 |

```
user@host> clear igmp statistics
```

```
user@host> show igmp statistics
```

```
IGMP packet statistics for all interfaces
```

| IGMP Message type | Received | Sent | Rx errors |
|-------------------------------------|----------|------|-----------|
| Membership Query | 0 | 0 | 0 |
| V1 Membership Report | 0 | 0 | 0 |
| DVMRP | 0 | 0 | 0 |
| PIM V1 | 0 | 0 | 0 |
| Cisco Trace | 0 | 0 | 0 |
| V2 Membership Report | 0 | 0 | 0 |
| Group Leave | 0 | 0 | 0 |
| Mtrace Response | 0 | 0 | 0 |
| Mtrace Request | 0 | 0 | 0 |
| Domain Wide Report | 0 | 0 | 0 |
| V3 Membership Report | 0 | 0 | 0 |
| Other Unknown types | | | 0 |
| IGMP v3 unsupported type | | | 0 |
| IGMP v3 source required for SSM | | | 0 |
| IGMP v3 mode not applicable for SSM | | | 0 |
| IGMP Global Statistics | | | |
| Bad Length | 0 | | |
| Bad Checksum | 0 | | |
| Bad Receive If | 0 | | |
| Rx non-local | 0 | | |

clear igmp-snooping membership

| | |
|---------------------------------|--|
| Syntax | <code>clear igmp-snooping membership</code> <code><vlan <i>vlan-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 9.1 for EX Series switches. |
| Description | Clear IGMP snooping dynamic membership information from the multicast forwarding table. |
| Options | none —Clear dynamic membership information for all VLANs. vlan <i>vlan-name</i> —(Optional) Clear dynamic membership information for the specified VLAN. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show igmp-snooping membership on page 220• clear igmp-snooping statistics on page 183 |
| List of Sample Output | clear igmp-snooping membership on page 182 |

Sample Output

clear igmp-snooping membership

```
user@switch> clear igmp-snooping membership vlan employee-vlan
```

clear igmp-snooping statistics

| | |
|---------------------------------|--|
| Syntax | clear igmp-snooping statistics |
| Release Information | Command introduced in Junos OS Release 9.1 for EX Series switches. |
| Description | Clear IGMP snooping statistics. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show igmp-snooping statistics on page 225• clear igmp-snooping membership on page 182 |
| List of Sample Output | clear igmp-snooping statistics on page 183 |

Sample Output

clear igmp-snooping statistics

```
user@switch> clear igmp-snooping statistics
```

clear mld-snooping membership

| | |
|---------------------------------|--|
| Syntax | clear mld-snooping membership <vlan <i>vlan-name</i>> |
| Release Information | Command introduced in Junos OS Release 12.1 for EX Series switches. |
| Description | Clear MLD snooping dynamic membership information from the multicast forwarding table. |
| Options | none —Clear dynamic membership information for all VLANs. vlan <i>vlan-name</i> —(Optional) Clear dynamic membership information for the specified VLAN. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show mld-snooping membership on page 230• clear mld-snooping statistics on page 185 |
| List of Sample Output | clear mld-snooping membership vlan employee-vlan on page 184 |

Sample Output

clear mld-snooping membership vlan employee-vlan

```
user@switch> clear mld-snooping membership vlan employee-vlan
```

clear mld-snooping statistics

| | |
|---------------------------------|--|
| Syntax | clear mld-snooping statistics |
| Release Information | Command introduced in Junos OS Release 12.1 for EX Series switches. |
| Description | Clear MLD snooping statistics. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show mld-snooping statistics on page 236• clear mld-snooping membership on page 184 |
| List of Sample Output | clear mld-snooping statistics on page 185 |

Sample Output

clear mld-snooping statistics

```
user@switch> clear mld-snooping statistics
```

clear multicast bandwidth-admission

| | |
|---------------------------------|---|
| Syntax | <pre>clear multicast bandwidth-admission <group <i>group-address</i>> <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <source <i>source-address</i>></pre> |
| Release Information | <p>Command introduced in Junos OS Release 8.3.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Reapply IP multicast bandwidth admissions. |
| Options | <p>none—Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.</p> <p>group <i>group-address</i>—(Optional) Reapply multicast bandwidth admissions for the specified group.</p> <p>inet—(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.</p> <p>inet6—(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.</p> <p>instance <i>instance-name</i>—(Optional) Reapply multicast bandwidth admission settings for the specified instance. If you do not specify an instance, the command applies to the master routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows:</p> <ul style="list-style-type: none">• If the interface is congested, and it was admitted previously, it is removed.• If the interface was rejected previously, the clear multicast bandwidth-admission command enables the interface to be admitted as long as enough bandwidth exists on the interface.• If you do not specify an interface, issuing the clear multicast bandwidth-admission command readmits any previously rejected interface for the relevant entries as long as enough bandwidth exists on the interface. <p>To manually reject previously admitted outbound interfaces, you must specify the interface.</p> <p>source <i>source-address</i>—(Optional) Use with the group option to reapply multicast bandwidth admission settings for the specified (source, group) entry.</p> |
| Required Privilege Level | clear |

- Related Documentation** • [show multicast interface on page 243](#)
- List of Sample Output** [clear multicast bandwidth-admission on page 187](#)
- Output Fields** When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear multicast bandwidth-admission](#)

```
user@host> clear multicast bandwidth-admission
```

clear multicast scope

| | |
|---|--|
| List of Syntax | Syntax on page 188 Syntax (EX Series Switch and the QFX Series) on page 188 |
| Syntax | <pre>clear multicast scope <inet inet6> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>clear multicast scope <inet inet6> <interface <i>interface-name</i>></pre> |
| Release Information | Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 option introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Clear IP multicast scope statistics. |
| Options | <p>none—(Same as logical-system all) Clear multicast scope statistics.</p> <p>inet—(Optional) Clear multicast scope statistics for IPv4 family addresses.</p> <p>inet6—(Optional) Clear multicast scope statistics for IPv6 family addresses.</p> <p>interface <i>interface-name</i>—(Optional) Clear multicast scope statistics on a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show multicast scope on page 265 |
| List of Sample Output | clear multicast scope on page 188 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear multicast scope

```
user@host> clear multicast scope
```

clear multicast sessions

| | |
|---|---|
| List of Syntax | Syntax on page 189 Syntax (EX Series Switch and the QFX Series) on page 189 |
| Syntax | clear multicast sessions <logical-system (all <i>logical-system-name</i>)> < <i>regular-expression</i> > |
| Syntax (EX Series Switch and the QFX Series) | clear multicast sessions < <i>regular-expression</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Clear IP multicast sessions. |
| Options | none —(Same as logical-system all) Clear multicast sessions. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>regular-expression</i> —(Optional) Clear only multicast sessions that contain the specified regular expression. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show multicast sessions on page 267 |
| List of Sample Output | clear multicast sessions on page 189 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear multicast sessions

```
user@host> clear multicast sessions
```

clear multicast statistics

| | |
|---|---|
| List of Syntax | Syntax on page 190 Syntax (EX Series Switch and the QFX Series) on page 190 |
| Syntax | <pre>clear multicast statistics <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>clear multicast statistics <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Clear IP multicast statistics. |
| Options | <p>none—Clear multicast statistics for all supported address families on all interfaces.</p> <p>inet—(Optional) Clear multicast statistics for IPv4 family addresses.</p> <p>inet6—(Optional) Clear multicast statistics for IPv6 family addresses.</p> <p>instance <i>instance-name</i>—(Optional) Clear multicast statistics for the specified instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear multicast statistics on a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show multicast statistics |
| List of Sample Output | clear multicast statistics on page 190 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear multicast statistics

```
user@host> clear multicast statistics
```

clear pim join

List of Syntax [Syntax on page 191](#)
 [Syntax \(EX Series Switch and the QFX Series\) on page 191](#)

Syntax clear pim join
 <group-address>
 <bidirectional | dense | sparse>
 <exact>
 <inet | inet6>
 <instance instance-name>
 <logical-system (all | logical-system-name)>
 <rp ip-address/prefix | source ip-address/prefix>
 <sg | star-g>

Syntax (EX Series Switch and the QFX Series) clear pim join
 <group-address>
 <dense | sparse>
 <exact>
 <inet | inet6>
 <instance instance-name>
 <rp ip-address/prefix | source ip-address/prefix>
 <sg | star-g>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
 Command introduced in Junos OS Release 11.3 for the QFX Series.
 Multiple new filter options introduced in Junos OS Release 13.2.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Clear the Protocol Independent Multicast (PIM) join and prune states.

Options **none**—Clear the PIM join and prune states for all groups, family addresses, and instances.

group-address—(Optional) Clear the PIM join and prune states for a group address.

bidirectional | dense | sparse—(Optional) Clear PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.

exact—(Optional) Clear only the group that exactly matches the specified group address.

inet | inet6—(Optional) Clear the PIM entries for IPv4 or IPv6 family addresses, respectively.

instance instance-name—(Optional) Clear the entries for a specific PIM-enabled routing instance.

logical-system (all | logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.

rp ip-address/prefix | source ip-address/prefix—(Optional) Clear the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

sg | star-g—(Optional) Clear PIM (S,G) or (*,G) entries.

Additional Information The **clear pim join** command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level clear

Related Documentation

- [show pim join on page 278](#)

List of Sample Output [clear pim join on page 192](#)
[clear pim join inet6 on page 192](#)
[clear pim join inet6 star-g on page 192](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pim join

```
user@host> clear pim join
Cleared 8 Join/Prune states
```

clear pim join inet6

```
user@host> clear pim join inet6
Cleared 4 Join/Prune states
```

clear pim join inet6 star-g

```
user@host> clear pim join inet6 star-g
Cleared 1 Join/Prune states
```

clear pim register

| | |
|---|---|
| List of Syntax | Syntax on page 193 Syntax (EX Series Switch and the QFX Series) on page 193 Syntax (PTX Series) on page 193 |
| Syntax | <pre>clear pim register <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>clear pim register <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>></pre> |
| Syntax (PTX Series) | <pre>clear pim register <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Release Information | <p>Command introduced in Junos OS Release 7.6.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Clear Protocol Independent Multicast (PIM) register message counters. |
| Options | <p>none—Clear PIM register message counters for all family addresses, instances, and interfaces.</p> <p>inet inet6—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Clear register message counters for a specific PIM-enabled routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear PIM register message counters for a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information | The clear pim register command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled. |
| Required Privilege Level | clear |

Related Documentation • [show pim statistics on page 314](#)

List of Sample Output [clear pim register on page 194](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear pim register](#)

```
user@host> clear pim register
```


clear pim statistics

| | |
|---|---|
| List of Syntax | Syntax on page 195 Syntax (EX Series Switch and the QFX Series) on page 195 |
| Syntax | <pre>clear pim statistics <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>clear pim statistics <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Clear Protocol Independent Multicast (PIM) statistics. |
| Options | <p>none—Clear PIM statistics for all family addresses, instances, and interfaces.</p> <p>inet inet6—(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Clear statistics for a specific PIM-enabled routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear PIM statistics for a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information | The clear pim statistics command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show pim statistics on page 314 |
| List of Sample Output | clear pim statistics on page 196 |
| Output Fields | See show pim statistics for an explanation of output fields. |

Sample Output

clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type      Received      Sent  Rx errors
Hello                  0             0      0
Register               0             0      0
Register Stop          0             0      0
Join Prune             0             0      0
Bootstrap              0             0      0
Assert                0             0      0
Graft                  0             0      0
Graft Ack              0             0      0
Candidate RP           0             0      0
V1 Query               2111          4222      0
V1 Register            0             0      0
V1 Register Stop       0             0      0
V1 Join Prune          14200         13115      0
V1 RP Reachability     0             0      0
V1 Assert              0             0      0
V1 Graft               0             0      0
V1 Graft Ack           0             0      0
PIM statistics summary for all interfaces:
Unknown type           0
V1 Unknown type        0
Unknown Version        0
Neighbor unknown       0
Bad Length             0
Bad Checksum           0
Bad Receive If         0
Rx Intf disabled       2007
Rx V1 Require V2       0
Rx Register not RP     0
RP Filtered Source     0
Unknown Reg Stop       0
Rx Join/Prune no state 1040
Rx Graft/Graft Ack no state 0
...
```

```
user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type      Received      Sent  Rx errors
Hello                  0             0      0
Register               0             0      0
Register Stop          0             0      0
Join Prune             0             0      0
Bootstrap              0             0      0
Assert                0             0      0
Graft                  0             0      0
Graft Ack              0             0      0
Candidate RP           0             0      0
V1 Query               1             0      0
V1 Register            0             0      0
...
```


mtrace

| | |
|---------------------------------|--|
| Syntax | <code>mtrace source</code> <code><logical-system logical-system-name></code> <code><routing-instance routing-instance-name></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 9.5 for SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 12.3 for the PTX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Display trace information about an IP multicast path. |
| Options | source —Source hostname or address. logical-system (logical-system-name) —(Optional) Perform this operation on a logical system. routing-instance routing-instance-name —(Optional) Trace a particular routing instance. |
| Additional Information | The mtrace command for multicast traffic is similar to the traceroute command used for unicast traffic. Unlike traceroute , mtrace traces traffic backwards, from the receiver to the source. |
| Required Privilege Level | view |
| List of Sample Output | mtrace source on page 200 |
| Output Fields | Table 7 on page 198 describes the output fields for the mtrace command. Output fields are listed in the approximate order in which they appear. |

Table 7: mtrace Output Fields

| Field Name | Field Description |
|-----------------------------------|---|
| Mtrace from | IP address of the receiver. |
| to | IP address of the source. |
| via group | IP address of the multicast group (if any). |
| Querying full reverse path | Indicates the full reverse path query has begun. |
| number-of-hops | Number of hops from the source to the named router or switch. |
| router-name | Name of the router or switch for this hop. |

Table 7: mtrace Output Fields (*continued*)

| Field Name | Field Description |
|-----------------|--|
| <i>address</i> | Address of the router or switch for this hop. |
| <i>protocol</i> | Protocol used (for example, PIM). |
| Round trip time | Average round-trip time, in milliseconds (ms). |
| total ttl of | Time-to-live (TTL) threshold. |

Sample Output

mtrace source

```
user@host> mtrace 192.1.4.2
Mtrace from 192.1.4.2 to 192.1.1.2 via group 0.0.0.0
Querying full reverse path... * *
 0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-2 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
-3 hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.
```

mtrace from-source

Syntax `mtrace from-source source source`
`<brief | detail>`
`<extra-hops extra-hops>`
`<group group>`
`<interval interval>`
`<loop>`
`<max-hops max-hops>`
`<max-queries max-queries>`
`<multicast-response | unicast-response>`
`<no-resolve>`
`<no-router-alert>`
`<response response>`
`<routing-instance routing-instance-name>`
`<tll tll>`
`<wait-time wait-time>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.3 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display trace information about an IP multicast path from a source to this router or switch. If you specify a group address with this command, Junos OS returns additional information, such as packet rates and losses.

Options **brief | detail**—(Optional) Display the specified level of output.

extra-hops *extra-hops*—(Optional) Number of hops to take after reaching a nonresponsive router. You can specify a number between **0** and **255**.

group *group*—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

interval *interval*—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10** seconds.

loop—(Optional) Loop indefinitely, displaying rate and loss statistics.

max-hops *max-hops*—(Optional) Maximum hops to trace toward the source. The range of values is **0** through **255**. The default value is **32** hops.

max-queries *max-queries*—(Optional) Maximum number of query attempts for any hop. The range of values is **1** through **32**. The default is **3**.

multicast-response—(Optional) Always request the response using multicast.

no-resolve—(Optional) Do not attempt to display addresses symbolically.

no-router-alert—(Optional) Do not use the router-alert IP option.

response *response*—(Optional) Send trace response to a host or multicast address.

routing-instance *routing-instance-name*—(Optional) Trace a particular routing instance.

source *source*—Source hostname or address.

ttl *tll*—(Optional) IP time-to-live (TTL) value. You can specify a number between 0 and 255. Local queries to the multicast group use a value of 1. Otherwise, the default value is 127.

unicast-response—(Optional) Always request the response using unicast.

wait-time *wait-time*—(Optional) Number of seconds to wait for a response. The default value is 3.

Required Privilege Level view

List of Sample Output [mtrace from-source on page 203](#)

Output Fields [Table 8 on page 202](#) describes the output fields for the **mtrace from-source** command. Output fields are listed in the approximate order in which they appear.

Table 8: mtrace from-source Output Fields

| Field Name | Field Description |
|-----------------------------------|---|
| Mtrace from | IP address of the receiver. |
| to | IP address of the source. |
| via group | IP address of the multicast group (if any). |
| Querying full reverse path | Indicates the full reverse path query has begun. |
| number-of-hops | Number of hops from the source to the named router or switch. |
| router-name | Name of the router or switch for this hop. |
| address | Address of the router or switch for this hop. |
| protocol | Protocol used (for example, PIM). |
| Round trip time | Average round-trip time, in milliseconds (ms). |
| total ttl of | Time-to-live (TTL) threshold. |
| source | Source address. |
| Response Dest | Response destination address. |
| Overall | Average packet rate for all traffic at each hop. |

Table 8: mtrace from-source Output Fields (*continued*)

| Field Name | Field Description |
|---|--|
| Packet Statistics for Traffic From | Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop. |
| Receiver | IP address receiving the multicast. |
| Query source | IP address sending the mtrace query. |

Sample Output

mtrace from-source

```

user@host> mtrace from-source source 192.1.4.2 group 225.1.1.1
Mtrace from 192.1.4.2 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
  0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-2 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
-3 hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.

Waiting to accumulate statistics...Results after 10 seconds:

Source      Response Dest    Overall    Packet Statistics For Traffic From
192.1.4.2 192.1.1.2 Packet    192.1.4.2 To 225.1.1.1
      v    ___/ rtt    2 ms    Rate    Lost/Sent = Pct Rate
192.1.2.1
192.1.3.2 routerC.lab.mycompany.net
      v    ^    ttl    2          0/0    = --    0 pps
192.1.4.1
192.1.2.2 routerB.lab.mycompany.net
      v    \__  ttl    3          ?/0          0 pps
192.1.1.2 192.1.1.2
Receiver    Query Source

```

mtrace monitor

| | |
|---------------------------------|---|
| Syntax | mtrace monitor |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Listen passively for IP multicast responses. To exit the mtrace monitor command, type Ctrl+c. |
| Options | none —Trace the master instance. |
| Required Privilege Level | view |
| List of Sample Output | mtrace monitor on page 205 |
| Output Fields | Table 9 on page 204 describes the output fields for the mtrace monitor command. Output fields are listed in the approximate order in which they appear. |

Table 9: mtrace monitor Output Fields

| Field Name | Field Description |
|-------------------------|---|
| Mtrace query at | Date and time of the query. |
| by | Address of the host issuing the query. |
| resp to | Response destination. |
| qid | Query ID number. |
| packet from...to | IP address of the query source and default group destination. |
| from...to | IP address of the multicast source and the response address. |
| via group | IP address of the group to trace. |
| mxhop | Maximum hop setting. |

Sample Output

mtrace monitor

```
user@host> mtrace monitor
Mtrace query at Oct 22 13:36:14 by 192.1.3.2, resp to 224.0.1.32, qid 74a5b8
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:17 by 192.1.3.2, resp to 224.0.1.32, qid 1d07ba
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:20 by 192.1.3.2, resp to same, qid 2fea1d
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:30 by 192.1.3.2, resp to same, qid 7c88ad
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)
```

mtrace to-gateway

Syntax `mtrace to-gateway gateway gateway`
`<brief | detail>`
`<extra-hops extra-hops>`
`<group group>`
`<interface interface-name>`
`<interval interval>`
`<loop>`
`<max-hops max-hops>`
`<max-queries max-queries>`
`<multicast-response | unicast-response>`
`<no-resolve>`
`<no-router-alert>`
`<response response>`
`<routing-instance routing-instance-name>`
`<tll ttl>`
`<unicast-response>`
`<wait-time wait-time>`

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display trace information about a multicast path from this router or switch to a gateway router or switch.

Options `gateway gateway`—Send the trace query to a gateway multicast address.

`brief | detail`—(Optional) Display the specified level of output.

`extra-hops extra-hops`—(Optional) Number of hops to take after reaching a nonresponsive router or switch. You can specify a number between **0** and **255**.

`group group`—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

`interface interface-name`—(Optional) Source address for sending the trace query.

`interval interval`—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10**.

`loop`—(Optional) Loop indefinitely, displaying rate and loss statistics.

`max-hops max-hops`—(Optional) Maximum hops to trace toward the source. You can specify a number between **0** and **255**. The default value is **32**.

`max-queries max-queries`—(Optional) Maximum number of query attempts for any hop. You can specify a number between **0** and **255**. The default value is **3**.

`multicast-response`—(Optional) Always request the response using multicast.

no-resolve—(Optional) Do not attempt to display addresses symbolically.

no-router-alert—(Optional) Do not use the router-alert IP option.

response *response*—(Optional) Send trace response to a host or multicast address.

routing-instance *routing-instance-name*—(Optional) Trace a particular routing instance.

ttl *tll*—(Optional) IP time-to-live value. You can specify a number between **0** and **225**.
Local queries to the multicast group use TTL 1. Otherwise, the default value is **127**.

unicast-response—(Optional) Always request the response using unicast.

wait-time *wait-time*—(Optional) Number of seconds to wait for a response. The default value is **3**.

Required Privilege Level view

List of Sample Output [mtrace to-gateway on page 207](#)

Output Fields [Table 10 on page 207](#) describes the output fields for the **mtrace to-gateway** command. Output fields are listed in the approximate order in which they appear.

Table 10: mtrace to-gateway Output Fields

| Field Name | Field Description |
|-----------------------------------|---|
| Mtrace from | IP address of the receiver. |
| to | IP address of the source. |
| via group | IP address of the multicast group (if any). |
| Querying full reverse path | Indicates the full reverse path query has begun. |
| <i>number-of-hops</i> | Number of hops from the source to the named router or switch. |
| <i>router-name</i> | Name of the router or switch for this hop. |
| <i>address</i> | Address of the router or switch for this hop. |
| <i>protocol</i> | Protocol used (for example, PIM). |
| Round trip time | Average round-trip time, in milliseconds (ms). |
| total ttl of | Time-to-live (TTL) threshold. |

Sample Output

mtrace to-gateway

```
user@host> mtrace to-gateway gateway 192.1.3.2 group 225.1.1.1 interface 192.1.1.73 brief
```

```
Mtrace from 192.1.1.73 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
 0  routerA.lab.mycompany.net (192.1.1.2)
-1  routerA.lab.mycompany.net (192.1.1.2) PIM thresh^ 1
-2  routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-3  routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
Round trip time 2 ms; total ttl of 3 required.
```

show igmp group

| | |
|---|---|
| List of Syntax | Syntax on page 209 Syntax (EX Series Switch and the QFX Series) on page 209 |
| Syntax | <pre>show igmp group <brief detail> <group-name> <logical-system (all logical-system-name)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show igmp group <brief detail> <group-name></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display Internet Group Management Protocol (IGMP) group membership information. |
| Options | <p>none—Display standard information about membership for all IGMP groups.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display group membership for the specified IP address only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear igmp membership on page 177 |
| List of Sample Output | show igmp group (Include Mode) on page 210 show igmp group (Exclude Mode) on page 211 show igmp group brief on page 211 show igmp group detail on page 211 |
| Output Fields | <p>Table 11 on page 209 describes the output fields for the show igmp group command. Output fields are listed in the approximate order in which they appear.</p> |

Table 11: show igmp group Output Fields

| Field Name | Field Description | Level of Output |
|------------|---|-----------------|
| Interface | Name of the interface that received the IGMP membership report. A name of local indicates that the local routing device joined the group itself. | All levels |
| Group | Group address. | All levels |

Table 11: show igmp group Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|---|-------------------|
| Group Mode | Mode the SSM group is operating in: Include or Exclude . | All levels |
| Source | Source address. | All levels |
| Source timeout | Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer. | detail |
| Last reported by | Address of the host that last reported membership in this group. | All levels |
| Timeout | Time remaining until the group membership is removed. | brief none |
| Group timeout | Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer. | detail |
| Type | Type of group membership: <ul style="list-style-type: none"> • Dynamic—Host reported the membership. • Static—Membership is configured. | All levels |

Sample Output

show igmp group (Include Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.2
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.3
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.4
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.2
    Group mode: Include
    Source: 10.0.0.4
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Source: 0.0.0.0

```



```

        Last reported by: Local
        Timeout:          0 Type: Dynamic
Group: 224.0.0.22
Source: 0.0.0.0
Last reported by: Local
Timeout:          0 Type: Dynamic

```

show igmp group (Exclude Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:          0 Type: Dynamic
  Group: 224.0.0.22
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:          0 Type: Dynamic

```

show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

show igmp group detail

```

user@host> show igmp group detail
Interface: t1-0/1/0.0
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.2
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.3
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.4
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
  Group: 232.1.1.2
    Group mode: Include
    Source: 10.0.0.4
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local

```

```
Group: 224.0.0.2
  Group mode: Exclude
  Source: 0.0.0.0
  Source timeout: 0
  Last reported by: Local
  Group timeout:      0 Type: Dynamic
Group: 224.0.0.22
  Group mode: Exclude
  Source: 0.0.0.0
  Source timeout: 0
  Last reported by: Local
  Group timeout:      0 Type: Dynamic
```

show igmp interface

| | |
|---|---|
| List of Syntax | Syntax on page 213 Syntax (EX Series Switches and the QFX Series) on page 213 |
| Syntax | <pre>show igmp interface <brief detail> <interface-name> <logical-system (all logical-system-name)></pre> |
| Syntax (EX Series Switches and the QFX Series) | <pre>show igmp interface <brief detail> <interface-name></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display information about Internet Group Management Protocol (IGMP)-enabled interfaces. |
| Options | <p>none—Display standard information about all IGMP-enabled interfaces.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface-name—(Optional) Display information about the specified IGMP-enabled interface only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear igmp membership on page 177 |
| List of Sample Output | show igmp interface on page 215 show igmp interface brief on page 216 show igmp interface detail on page 216 show igmp interface <interface-name> on page 216 |
| Output Fields | <p>Table 12 on page 213 describes the output fields for the show igmp interface command. Output fields are listed in the approximate order in which they appear.</p> |

Table 12: show igmp interface Output Fields

| Field Name | Field Description | Level of Output |
|------------|------------------------|-----------------|
| Interface | Name of the interface. | All levels |

Table 12: show igmp interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|---|-----------------|
| Querier | Address of the routing device that has been elected to send membership queries. | All levels |
| State | State of the interface: Up or Down . | All levels |
| SSM Map Policy | Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface. | All levels |
| Timeout | How long until the IGMP querier is declared to be unreachable, in seconds. | All levels |
| Version | IGMP version being used on the interface: 1, 2, or 3. | All levels |
| Groups | Number of groups on the interface. | All levels |
| Group limit | Maximum number of groups allowed on the interface. Any joins requested after the limit is reached are rejected. | All levels |
| Group threshold | Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message. | All levels |
| Group log-interval | Time (in seconds) between consecutive log messages. | All levels |
| Immediate Leave | State of the immediate leave option: <ul style="list-style-type: none"> On—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface. Off—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds. | All levels |
| Promiscuous Mode | State of the promiscuous mode option: <ul style="list-style-type: none"> On—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces. Off—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces. | All levels |

Table 12: show igmp interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------|--|-----------------|
| Passive | <p>State of the passive mode option:</p> <ul style="list-style-type: none"> • On—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves. • Off—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves. <p>The passive statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the on state declaration:</p> <ul style="list-style-type: none"> • send-general-query—The interface sends general queries. • send-group-query—The interface sends group-specific and group-source-specific queries. • allow-receive—The interface receives control traffic. | All levels |
| OIF map | Name of the OIF map (if configured) associated with the interface. | All levels |
| SSM map | Name of the source-specific multicast (SSM) map (if configured) used on the interface. | All levels |
| Configured Parameters | <p>Information configured by the user:</p> <ul style="list-style-type: none"> • IGMP Query Interval—Interval (in seconds) at which this router sends membership queries when it is the querier. • IGMP Query Response Interval—Time (in seconds) that the router waits for a report in response to a general query. • IGMP Last Member Query Interval—Time (in seconds) that the router waits for a report in response to a group-specific query. • IGMP Robustness Count—Number of times the router retries a query. | All levels |
| Derived Parameters | <p>Derived information:</p> <ul style="list-style-type: none"> • IGMP Membership Timeout—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed. • IGMP Other Querier Present Timeout—Time (in seconds) that the router waits for the IGMP querier to send a query. | All levels |

Sample Output

show igmp interface

```

user@host> show igmp interface
Interface: at-0/3/1.0
  Querier: 10.111.30.1
  State:      Up Timeout:  None Version:  2 Groups:    4
  SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
  Querier: 10.111.10.1
  State:      Up Timeout:  None Version:  2 Groups:    2
  SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
  Querier: 10.111.20.1

```

```
State:          Up Timeout:    None Version:  2 Groups:      4
SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 215](#).

show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 215](#).

show igmp interface <interface-name>

```
user@host# show igmp interface ge-3/2/0.0
Interface: ge-3/2/0.0
Querier: 20.1.1.1
State: Up Timeout:    None Version:  3 Groups:      1
Group limit: 8
Group threshold: 60
Group log-interval: 10
Immediate leave: Off
Promiscuous mode: Off
```

show igmp statistics

| | |
|---|--|
| List of Syntax | Syntax on page 217 Syntax (EX Series Switch and the QFX Series) on page 217 |
| Syntax | <pre>show igmp statistics <brief detail> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show igmp statistics <brief detail> <interface <i>interface-name</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display Internet Group Management Protocol (IGMP) statistics. |
| Options | <p>none—Display IGMP statistics for all interfaces.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display IGMP statistics about the specified interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear igmp statistics on page 180 |
| List of Sample Output | show igmp statistics on page 218 show igmp statistics interface on page 219 |
| Output Fields | <p>Table 13 on page 217 describes the output fields for the show igmp statistics command. Output fields are listed in the approximate order in which they appear.</p> |

Table 13: show igmp statistics Output Fields

| Field Name | Field Description |
|------------------------|--|
| IGMP packet statistics | Heading for IGMP packet statistics for all interfaces or for the specified interface name. |

Table 13: show igmp statistics Output Fields (*continued*)

| Field Name | Field Description |
|------------------------|---|
| IGMP Message type | <p>Summary of IGMP statistics:</p> <ul style="list-style-type: none"> • Membership Query—Number of membership queries sent and received. • V1 Membership Report—Number of version 1 membership reports sent and received. • DVMRP—Number of DVMRP messages sent or received. • PIM V1—Number of PIM version 1 messages sent or received. • Cisco Trace—Number of Cisco trace messages sent or received. • V2 Membership Report—Number of version 2 membership reports sent or received. • Group Leave—Number of group leave messages sent or received. • Mtrace Response—Number of Mtrace response messages sent or received. • Mtrace Request—Number of Mtrace request messages sent or received. • Domain Wide Report—Number of domain-wide reports sent or received. • V3 Membership Report—Number of version 3 membership reports sent or received. • Other Unknown types—Number of unknown message types received. • IGMP v3 unsupported type—Number of messages received with unknown and unsupported IGMP version 3 message types. • IGMP v3 source required for SSM—Number of IGMP version 3 messages received that contained no source. • IGMP v3 mode not applicable for SSM—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM). |
| Received | Number of messages received. |
| Sent | Number of messages sent. |
| Rx errors | Number of received packets that contained errors. |
| IGMP Global Statistics | <p>Summary of IGMP statistics for all interfaces.</p> <ul style="list-style-type: none"> • Bad Length—Number of messages received with length errors so severe that further classification could not occur. • Bad Checksum—Number of messages received with a bad IP checksum. No further classification was performed. • Bad Receive If—Number of messages received on an interface not enabled for IGMP. • Rx non-local—Number of messages received from senders that are not local. • Timed out—Number of groups that timed out as a result of not receiving an explicit leave message. • Rejected Report—Number of reports dropped because of the IGMP group policy. • Total Interfaces—Number of interfaces configured to support IGMP. |

Sample Output

show igmp statistics

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        8883         459      0
V1 Membership Report     0            0      0

```


| | | | |
|-------------------------------------|------|---|---|
| DVMRP | 0 | 0 | 0 |
| PIM V1 | 0 | 0 | 0 |
| Cisco Trace | 0 | 0 | 0 |
| V2 Membership Report | 0 | 0 | 0 |
| Group Leave | 0 | 0 | 0 |
| Mtrace Response | 0 | 0 | 0 |
| Mtrace Request | 0 | 0 | 0 |
| Domain Wide Report | 0 | 0 | 0 |
| V3 Membership Report | 0 | 0 | 0 |
| Other Unknown types | | | 0 |
| IGMP v3 unsupported type | | | 0 |
| IGMP v3 source required for SSM | | | 0 |
| IGMP v3 mode not applicable for SSM | | | 0 |
| IGMP Global Statistics | | | |
| Bad Length | 0 | | |
| Bad Checksum | 0 | | |
| Bad Receive If | 0 | | |
| Rx non-local | 1227 | | |
| Timed out | 0 | | |
| Rejected Report | 0 | | |
| Total Interfaces | 2 | | |

show igmp statistics interface

```

user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type      Received      Sent  Rx errors
Membership Query        0           230      0
V1 Membership Report    0           0        0

```

show igmp-snooping membership

| | |
|---------------------------------|--|
| Syntax | <pre>show igmp-snooping membership <brief detail> <interface <i>interface-name</i>> <vlan (<i>vlan-id</i> <i>vlan-name</i>)></pre> |
| Release Information | Command introduced in Junos OS Release 9.1 for EX Series switches. |
| Description | Display the multicast group membership information maintained by IGMP snooping. |
| Options | <p>none—Display the multicast group membership information about all VLANs on which IGMP snooping is enabled.</p> <p>brief detail—(Optional) Display the specified level of output. The default is brief.</p> <p>interface <i>interface-name</i>—(Optional) Display the multicast group membership information about the specified interface.</p> <p>vlan (<i>vlan-id</i> <i>vlan-name</i>)—(Optional) Display the multicast group membership for the specified VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show igmp-snooping route on page 223 • show igmp-snooping statistics on page 225 • show igmp-snooping vlans on page 227 • Verifying IGMP Snooping (CLI Procedure) on page 168 • Configuring IGMP Snooping (CLI Procedure) on page 39 |
| List of Sample Output | <p>show igmp-snooping membership on page 221</p> <p>show igmp-snooping membership detail on page 221</p> <p>show igmp-snooping membership vlan detail on page 222</p> |
| Output Fields | Table 14 on page 220 lists the output fields for the show igmp-snooping membership command. Output fields are listed in the approximate order in which they appear. |

Table 14: show igmp-snooping membership Output Fields

| Field Name | Field Description | Level of Output |
|------------|--|-----------------|
| VLAN | Name of the VLAN. | All |
| Interfaces | Interfaces that are members of the listed multicast group. | All |
| Tag | Numerical identifier of the VLAN. | detail |

Table 14: show igmp-snooping membership Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------|---|-----------------|
| Router interfaces | <p>List of information about multicast-router interfaces:</p> <ul style="list-style-type: none"> Name of the multicast-router interface. static or dynamic—Whether the multicast-router interface is static or dynamic. Uptime—For static interfaces, amount of time since the interface was configured as a multicast-router interface or since the interface last flapped. For dynamic interfaces, amount of time since the first query was received on the interface or since the interface last flapped. timeout—Seconds remaining before a dynamic multicast-router interface times out. | detail |
| Group | <p>IP multicast address of the multicast group.</p> <p>The following information is provided for the multicast group:</p> <ul style="list-style-type: none"> Name of the interface belonging to the multicast group. Last reporter—Last host to report membership for the multicast group. Receiver count—Number of hosts on the interface that are members of the multicast group. This field appears only if immediate-leave is configured on the VLAN. Uptime—Length of time (in hours, minutes, and seconds) a multicast group has been active on the interface. timeout—Time (in seconds) left until the entry for the multicast group is removed from the multicast group if no membership reports are received on the interface. This counter is reset to its maximum value when a membership report is received. Flags—The lowest IGMP version in use by a host that is a member of the group on the interface. If the flag static is included, the interface has been configured as static member of the multicast group. Include source—Multicast source addresses of all IGMPv3 membership reports received for the group on the interface. | detail |

Sample Output

show igmp-snooping membership

```

user@switch> show igmp-snooping membership
VLAN: vlan24
  224.1.1.1      *
    Interfaces: ge-0/0/0.0
  224.1.1.100    *
    Interfaces: ge-0/0/0.0
  225.1.1.100    *
    Interfaces: ge-0/0/0.0

```

show igmp-snooping membership detail

```

user@switch> show igmp-snooping membership detail

VLAN: vlan2 Tag: 2 (Index: 3)
Router interfaces:

```

```

    ge-1/0/0.0 dynamic Uptime: 00:14:24 timeout: 253
Group: 225.0.0.1
    ge-1/0/17.0 259 Last reporter: 13.0.0.90 Receiver count: 1
    Uptime: 00:00:19 timeout: 259 Flags: <V3-hosts>
    Include source: 10.2.11.5, 10.2.11.12

```

show igmp-snooping membership vlan detail

```

user@switch> show igmp-snooping membership vlan vlan700 detail
VLAN: vlan700 Tag: 700 (Index: 52)
  Router interfaces:
    ae2.0 dynamic Uptime: 16:53:13 timeout: 245
  Group: 230.150.10.1
    ge-0/0/1.0 Last reporter: 100.2.188.201
    Uptime: 17:00:52 timeout: 237 Flags: <V2-hosts>
    ge-0/0/0.0 Last reporter: 100.2.188.202
    Uptime: 17:00:50 timeout: 243 Flags: <V2-hosts>

```

show igmp-snooping route

| | |
|---------------------------------|---|
| Syntax | <code>show igmp-snooping route</code> <code><brief detail></code> <code><ethernet-switching inet></code> <code><vlan (vlan-id vlan-name)></code> |
| Release Information | Command introduced in Junos OS Release 9.1 for EX Series switches. |
| Description | Display IGMP snooping route information. |
| Options | <p>none—Display route information for all VLANs on which IGMP snooping is enabled.</p> <p>brief detail—(Optional) Display the specified level of output. The default is brief.</p> <p>ethernet-switching—(Optional) Display information on Layer 2 multicast routes. This is the default.</p> <p>inet—(Optional) Display information for Layer 3 multicast routes.</p> <p>vlan (vlan-id vlan-name)—(Optional) Display route information for the specified VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show igmp-snooping membership on page 220 • show igmp-snooping statistics on page 225 • show igmp-snooping vlans on page 227 • Verifying IGMP Snooping (CLI Procedure) on page 168 • Configuring IGMP Snooping (CLI Procedure) on page 39 |
| List of Sample Output | show igmp-snooping route vlan v18 on page 224 show igmp-snooping route detail on page 224 show igmp-snooping route inet detail on page 224 |
| Output Fields | Table 15 on page 223 lists the output fields for the show igmp-snooping route command. Output fields are listed in the approximate order in which they appear. |

Table 15: show igmp-snooping route Output Fields

| Field Name | Field Description |
|---------------|---|
| Table | Routing table ID for virtual routing instances. |
| Routing Table | Routing table ID for virtual routing instances. |
| VLAN | Name of the VLAN on which IGMP snooping is enabled. |
| Group | Multicast IPv4 group address. |

Table 15: show igmp-snooping route Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|--|
| Next-hop | ID associated with the next-hop device. |
| Routing next-hop | ID associated with the Layer 3 next-hop device. |
| Interface or Interfaces | Name of the interface or interfaces in the VLAN associated with the multicast group. |
| Layer 2 next-hop | ID associated with the Layer 2 next-hop device. |

Sample Output

show igmp-snooping route vlan v18

```

user@switch> show igmp-snooping route vlan v18
VLAN      Group      Next-hop
v1an18    224.0.0.0, *
v1an18    225.20.20.1, *    1539

```

show igmp-snooping route detail

```

user@switch> show igmp-snooping route detail
VLAN      Group      Next-hop
default   224.0.0.0, *
v1an100    224.0.0.0, *    1332
          Interfaces: ge-1/0/1.0
VLAN      Group      Next-hop
v1an100    226.0.0.1, *    1334
          Interfaces: ge-1/0/1.0, ge-5/0/30.0

```

show igmp-snooping route inet detail

```

user@switch> show igmp-snooping route inet detail
Routing table: 0
Group: 229.0.0.1, 171.2.60.100
  Routing next-hop: 3448
  vlan.100
    Interface: vlan.100, VLAN: vlan100, Layer 2 next-hop: 3343

```

show igmp-snooping statistics

| | |
|---------------------------------|--|
| Syntax | show igmp-snooping statistics |
| Release Information | Command introduced in Junos OS Release 9.1 for EX Series switches. |
| Description | Display IGMP snooping statistics. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear igmp-snooping statistics on page 183 • show igmp-snooping membership on page 220 • show igmp-snooping route on page 223 • show igmp-snooping vlans on page 227 • Verifying IGMP Snooping (CLI Procedure) on page 168 • Configuring IGMP Snooping (CLI Procedure) on page 39 |
| List of Sample Output | show igmp-snooping statistics on page 226 |
| Output Fields | Table 16 on page 225 lists the output fields for the show igmp-snooping statistics command. Output fields are listed in the approximate order in which they appear. |

Table 16: show igmp-snooping statistics Output Fields

| Field Name | Field Description |
|-------------------|--|
| Bad length | IGMP packet has illegal or bad length. |
| Bad checksum | IGMP or IP checksum is incorrect. |
| Invalid interface | Packet was received through an invalid interface. |
| Not local | Not used—always 0. |
| Receive unknown | Unknown IGMP type. |
| Timed out | Not used—always 0. |
| IGMP Type | Type of IGMP message (Query, Report, Leave, or Other). |
| Received | Number of IGMP packets received. |
| Transmitted | Number of IGMP packets transmitted. |
| Recv Errors | Number of packets received that did not conform to the IGMP version 1 (IGMPv1), IGMPv2, or IGMPv3 standards. |

Sample Output

show igmp-snooping statistics

```
user@switch> show igmp-snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 0
```

| IGMP Type | Received | Transmitted | Recv Errors |
|-----------|----------|-------------|-------------|
| Queries: | 74295 | 0 | 0 |
| Reports: | 18148423 | 0 | 16333523 |
| Leaves: | 0 | 0 | 0 |
| Other: | 0 | 0 | 0 |

show igmp-snooping vlans

| | |
|---------------------------------|---|
| Syntax | show igmp-snooping vlans <brief detail> <vlan (vlan-id vlan-name)> |
| Release Information | Command introduced in Junos OS Release 9.1 for EX Series switches. |
| Description | Display IGMP snooping information for a VLAN or for all VLANs. |
| Options | <p>none—Display IGMP snooping information for all VLANs on which IGMP snooping is enabled.</p> <p>brief detail—(Optional) Display the specified level of output. The default is brief.</p> <p>vlan (vlan-id vlan vlan-number)—(Optional) Display IGMP snooping information for the specified VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show igmp-snooping membership on page 220 • show igmp-snooping route on page 223 • show igmp-snooping statistics on page 225 • Verifying IGMP Snooping (CLI Procedure) on page 168 • Configuring IGMP Snooping (CLI Procedure) on page 39 |
| List of Sample Output | show igmp-snooping vlans on page 228 show igmp-snooping vlans vlan v10 on page 228 show igmp-snooping vlans detail on page 228 |
| Output Fields | Table 17 on page 227 lists the output fields for the show igmp-snooping vlans command. Output fields are listed in the approximate order in which they appear. |

Table 17: show igmp-snooping vlans Output Fields

| Field Name | Field Description | Level of Output |
|------------|--|-----------------|
| VLAN | Name of the VLAN. | All levels |
| Interfaces | Number of interfaces in the VLAN. | brief |
| Groups | Number of groups in the VLAN. | brief |
| MRouters | Number of multicast routers in the VLAN. | brief |
| Receivers | Number of VLAN interfaces with a receiver for any group. Indicates how many VLAN interfaces would receive data because of IGMP membership. | brief |

Table 17: show igmp-snooping vlans Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------|---|-----------------|
| RxVlans | Number of MVR receiver VLANs configured for that MVR source VLAN. | brief |
| Tag | VLAN tag. | detail |
| Membership timeout | The group membership timeout value, which determines how long the switch waits before removing an IGMP snooping group from its membership table if no report is received. | detail |
| Querier timeout | The maximum length of time the switch waits to take over as IGMP querier if no query is received. | detail |
| vlan-interface | The Layer 3 interface, if any, associated with the VLAN. | detail |
| Interface | <p>Name of the interface.</p> <p>The following information is provided for each interface:</p> <ul style="list-style-type: none"> tagged or untagged—Whether the interface accepts tagged packets (trunk mode and tagged-access mode ports) or untagged packets (access mode ports) Groups—The number of multicast groups the interface belongs to. Reporters—The number of hosts on the interface that are current members of multicast groups. This field appears only when <i>immediate-leave</i> is configured on the VLAN. Router—Indicates the interface is a multicast-router interface. | detail |

Sample Output

show igmp-snooping vlans

```

user@switch> show igmp-snooping vlans
VLAN          Interfaces Groups MRouters Receivers RxVlans
default              0      0      0         0        0
v1                  4      0      1         0        0
v10                 1      0      0         0        0
v11                 1      0      0         0        0
v180                3      0      1         0        0
v181                3      0      0         0        0
v182                3      0      0         0        0

```

show igmp-snooping vlans vlan v10

```

user@switch> show igmp-snooping vlans vlan v10
VLAN          Interfaces Groups MRouters Receivers RxVlans
v10              1      0      0         0        0

```

show igmp-snooping vlans detail

```

user@switch> show igmp-snooping vlans detail

VLAN: default, Tag: 0
Membership timeout: 54, Querier timeout: 52
VLAN: v2146-API, Tag: 2146, vlan-interface: vlan.2146

```

```
Membership timeout: 54, Querier timeout: 52
Interface: ae0.0, tagged, Groups: 0, Reporters: 0,
Interface: ge-7/0/21.0, untagged, Groups: 0, Reporters: 0
Interface: ge-1/0/24.0, untagged, Groups: 0, Reporters: 0
Interface: ge-1/0/25.0, untagged, Groups: 0, Reporters: 0
Interface: ge-1/0/26.0, untagged, Groups: 0, Reporters: 0
Interface: ge-1/0/36.0, untagged, Groups: 0, Reporters: 0
Interface: ge-1/0/37.0, untagged, Groups: 0, Reporters: 0
Interface: ge-1/0/38.0, untagged, Groups: 0, Reporters: 0
```

show mld-snooping membership

| | |
|---------------------------------|---|
| Syntax | <pre>show mld-snooping membership <brief detail> <interface <i>logical-interface-name</i>> <vlan (<i>vlan-id</i> <i>vlan-name</i>) ></pre> |
| Release Information | Command introduced in Junos OS Release 12.1 for EX Series switches. |
| Description | Display the multicast group membership information maintained by MLD snooping. |
| Options | <p>none—Display the multicast group membership information for all VLANs on which MLD snooping is enabled.</p> <p>brief detail—(Optional) Display the specified level of output. The default is brief.</p> <p>interface <i>interface-name</i>—(Optional) Display the multicast group membership information for the specified interface.</p> <p>vlan (<i>vlan-id</i> <i>vlan-name</i>)—(Optional) Display the multicast group membership for the specified VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear mld-snooping membership on page 184 • show mld-snooping route on page 233 • show mld-snooping statistics on page 236 • show mld-snooping vlans on page 238 • Verifying MLD Snooping (CLI Procedure) on page 170 • Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |
| List of Sample Output | show mld-snooping membership on page 231 show mld-snooping membership detail on page 232 |
| Output Fields | Table 18 on page 230 lists the output fields for the show mld-snooping membership command. Output fields are listed in the approximate order in which they appear. |

Table 18: show mld-snooping membership Output Fields

| Field Name | Field Description | Level of Output |
|-------------------|--|-----------------|
| VLAN | Name of the VLAN. | All |
| Interfaces | Interfaces that are members of the listed multicast group. | brief |
| Tag | Numerical identifier of the VLAN. | detail |

Table 18: show mld-snooping membership Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------|--|-----------------|
| Router interfaces | <p>List of information about multicast-router interfaces:</p> <ul style="list-style-type: none"> Name of the multicast-router interface. static or dynamic—Whether the multicast-router interface has been statically configured or dynamically learned. Uptime—For static interfaces, amount of time since the interface was configured as a multicast-router interface or since the interface last flapped. For dynamic interfaces, amount of time since the first query was received on the interface or since the interface last flapped. timeout—Seconds remaining before a dynamic multicast-router interface times out. | detail |
| Group | <p>IP multicast address of the multicast group.</p> <p>The following information is provided for the multicast group:</p> <ul style="list-style-type: none"> Name of the interface belonging to the multicast group. Timeout—Time (in seconds) left until a dynamically learned interface is removed from the multicast group if no MLD membership reports are received on the interface. This counter is reset to its maximum value when a membership report is received. Flags—The lowest MLD version in use by a host that is a member of the group on the interface. If the flag static is included, the interface has been configured as static member of the multicast group. Receiver count—Number of hosts on the interface that are members of the multicast group. This field appears only if immediate-leave is configured on the VLAN. Last reporter—Last host to report membership for the multicast group. Include source—Multicast source addresses from all MLDv2 membership reports received for the group on the interface. | detail |

Sample Output

show mld-snooping membership

```

user@switch> show mld-snooping membership
VLAN: mld_vlan
ff1e::2010
    Interfaces: ge-1/0/30.0
ff1e::2011
    Interfaces: ge-1/0/30.0
ff1e::2012
    Interfaces: ge-1/0/30.0
ff1e::2013
    Interfaces: ge-1/0/30.0

```

```
ff1e::2014
  Interfaces: ge-1/0/30.0
```

show mld-snooping membership detail

```
user@switch> show mld-snooping membership detail
VLAN: mld-vlan Tag: 100 (Index: 3)
  Router interfaces:
    ge-1/0/0.0 static Uptime: 00:57:13
  Group: ff1e::2010
    ge-1/0/30.0 Timeout: 180 Flags: <V2-hosts>
    Last reporter: fe80::2020:1:1:3
    Include source: 2020:1:1:1::2
VLAN: mld-vlan1 Tag: 200 (Index: 4)
  Router interfaces:
    ae200.0 dynamic Uptime: 00:14:24 timeout: 244
  Group: ff1e::2010
    ge-12/0/31.0 Timeout: 224 Flags: <V1-hosts>
    Last reporter: fe80::2020:1:1:4
```

show mld-snooping route

| | |
|---------------------------------|---|
| Syntax | <pre>show mld-snooping route <brief detail> <ethernet-switching inet6> <vlan (vlan-id vlan-name)></pre> |
| Release Information | Command introduced in Junos OS Release 12.1 for EX Series switches. |
| Description | Display multicast route information maintained by MLD snooping. |
| Options | <p>none—Display route information for all VLANs on which MLD snooping is enabled.</p> <p>brief detail—(Optional) Display the specified level of output. The default is brief.</p> <p>ethernet-switching—(Optional) Display information on Layer 2 IPv6 multicast routes. This is the default.</p> <p>inet6—(Optional) Display information on Layer 3 IPv6 multicast routes.</p> <p>vlan (vlan-id vlan-name) —(Optional) Display route information for the specified VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show mld-snooping membership on page 230 • show mld-snooping statistics on page 236 • show mld-snooping vlans on page 238 • Verifying MLD Snooping (CLI Procedure) on page 170 • Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |
| List of Sample Output | <p>show mld-snooping route on page 234</p> <p>show mld-snooping route detail on page 234</p> <p>show mld-snooping route inet6 detail on page 235</p> |
| Output Fields | Table 19 on page 233 lists the output fields for the show mld-snooping route command. Output fields are listed in the approximate order in which they appear. |

Table 19: show mld-snooping route Output Fields

| Field Name | Field Description |
|---------------|--|
| Table | Routing table ID for virtual routing instances. |
| Routing Table | Routing table ID for virtual routing instances. |
| VLAN | Name of the VLAN on which MLD snooping is enabled. |

Table 19: show mld-snooping route Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|--|
| Group | Multicast IPv6 group address. Only the last 32 bits of the address are shown. The switch uses only these bits in determining multicast routes. |
| Next-hop | ID associated with the next-hop device. |
| Routing next-hop | ID associated with the Layer 3 next-hop device. |
| Interface or Interfaces | Name of the interface or interfaces in the VLAN associated with the multicast group. |
| Layer 2 next-hop | ID associated with the Layer 2 next-hop device. |

Sample Output

show mld-snooping route

```

user@switch> show mld-snooping route

VLAN      Group      Next-hop
vlan1     ::0000:0001 1464
vlan1     ff00::
vlan10    ::0000:0002 1599
vlan10    ff00::
vlan11    ::0000:0002 1513
vlan11    ff00::
vlan12    ff00::
vlan13    ff00::
vlan14    ff00::
vlan15    ff00::
vlan16    ff00::
vlan17    ff00::
vlan18    ff00::
vlan19    ff00::
vlan2     ff00::
vlan20    ::0000:0002 1602
vlan20    ff00::
vlan3     ff00::
vlan4     ff00::
vlan5     ff00::
vlan6     ff00::
vlan7     ff00::
vlan8     ff00::
vlan9     ff00::
default  ff00::

```

show mld-snooping route detail

```

user@switch> show mld-snooping route detail
VLAN      Group      Next-hop
mld-vlan  ::0000:2010 1323
          Interfaces: ge-1/0/30.0
VLAN      Group      Next-hop
mld-vlan  ff00::      1317

```



```

          Interfaces: ge-1/0/0.0
VLAN      Group      Next-hop
mld-vlan  ::0000:0000  1317
          Interfaces: ge-1/0/0.0
VLAN      Group      Next-hop
mld-vlan1 ::0000:2010   1324
          Interfaces: ge-12/0/31.0
VLAN      Group      Next-hop
mld-vlan1 ff00::       1318
          Interfaces: ae200.0
VLAN      Group      Next-hop
mld-vlan1 ::0000:0000  1318
          Interfaces: ae200.0

```

show mld-snooping route inet6 detail

```

user@switch> show mld-snooping route inet6 detail
Routing table: 0
Group: ff05::1, 4001::11
  Routing next-hop: 1352
  vlan.2
    Interface: vlan.2, VLAN: vlan2, Layer 2 next-hop: 1387

```

show mld-snooping statistics

| | |
|---------------------------------|---|
| Syntax | <code>show mld-snooping statistics</code> |
| Release Information | Command introduced in Junos OS Release 12.1 for EX Series switches. |
| Description | Display MLD snooping statistics. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear mld-snooping statistics on page 185 • show mld-snooping membership on page 230 • show mld-snooping route on page 233 • show mld-snooping vlans on page 238 • Verifying MLD Snooping (CLI Procedure) on page 170 |
| List of Sample Output | show mld-snooping statistics on page 237 |
| Output Fields | <p>Table 20 on page 236 lists the output fields for the show mld-snooping statistics command. Output fields are listed in the approximate order in which they appear.</p> |

Table 20: show mld-snooping statistics Output Fields

| Field Name | Field Description |
|-------------------|--|
| Bad length | MLD packet has illegal or bad length. |
| Bad checksum | MLD or IP checksum is incorrect. |
| Invalid interface | Packet was received through an invalid interface. |
| Not Local | Not used—always 0. |
| Receive unknown | Unknown MLD message type. |
| Timed out | Not used—always 0. |
| MLD Type | Type of MLD message (Query, Report, Leaves, or Other). |
| Received | Number of MLD packets received. |
| Transmitted | Number of MLD packets transmitted. |
| Recv Errors | Number of packets received that did not conform to the MLD version 1 (MLDv1) or MLDv2 standards. |

Sample Output

show mld-snooping statistics

```
user@switch> show mld-snooping statistics
```

```
Bad length: 0 Bad checksum: 0 Invalid interface: 0
```

```
Not local: 0 Receive unknown: 0 Timed out: 0
```

| MLD Type | Received | Transmitted | Recv Errors |
|----------|----------|-------------|-------------|
| Queries: | 74295 | 0 | 0 |
| Reports: | 18148423 | 0 | 16333523 |
| Leaves: | 0 | 0 | 0 |
| Other: | 0 | 0 | 0 |

show mld-snooping vlans

| | |
|---------------------------------|--|
| Syntax | show mld-snooping vlans <brief detail> <vlan <i>vlan-name</i>> |
| Release Information | Command introduced in Junos OS Release 12.1 for EX Series switches. |
| Description | Display MLD snooping information for a VLAN or for all VLANs. |
| Options | <p>none—Display MLD snooping information for all VLANs on which MLD snooping is enabled.</p> <p>brief detail—(Optional) Display the specified level of output. The default is brief.</p> <p>vlan <i>vlan-name</i> —(Optional) Display MLD snooping information for the specified VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show mld-snooping membership on page 230 • show mld-snooping route on page 233 • show mld-snooping statistics on page 236 • Verifying MLD Snooping (CLI Procedure) on page 170 • Configuring MLD Snooping on a VLAN (CLI Procedure) on page 50 |
| List of Sample Output | <p>show mld-snooping vlans on page 239</p> <p>show mld-snooping vlans vlan v10 on page 239</p> <p>show mld-snooping vlans vlan vlan2 detail on page 239</p> <p>show mld-snooping vlans detail on page 239</p> |
| Output Fields | Table 21 on page 238 lists the output fields for the show mld-snooping vlans command. Output fields are listed in the approximate order in which they appear. |

Table 21: show mld-snooping vlans Output Fields

| Field Name | Field Description | Level of Output |
|------------|---|-----------------|
| VLAN | Name of the VLAN. | All levels |
| Interfaces | Number of interfaces in the VLAN. | brief |
| Groups | Number of groups in the VLAN. | brief |
| MRouters | Number of multicast-router interfaces in the VLAN. | brief |
| Receivers | Number of interfaces in the VLAN with a receiver for any group. Indicates how many interfaces might receive data because of MLD group membership. | brief |
| Tag | VLAN tag. | detail |

Table 21: show mld-snooping vlans Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------|---|-----------------|
| vlan-interface | The Layer 3 interface, if any, associated with the VLAN. | detail |
| Interface | <p>Name of the interface.</p> <p>The following information is provided for each interface:</p> <ul style="list-style-type: none"> • tagged or untagged—Whether the interface accepts tagged packets (trunk mode and tagged-access mode ports) or untagged packets (access mode ports) • Groups—The number of multicast groups the interface belongs to • Reporters—The number of hosts on the interface that are current members of multicast groups. This field appears only when immediate-leave is configured on the VLAN. • Router—Indicates the interface is a multicast-router interface | detail |

Sample Output

show mld-snooping vlans

```

user@switch> show mld-snooping vlans
VLAN          Interfaces Groups MRouters Receivers
default              0      0      0      0
v1                 11     50      0      0
v10                 1      0      0      0
v11                 1      0      0      0
v180                 3      0      1      0
v181                 3      0      0      0
v182                 3      0      0      0

```

show mld-snooping vlans vlan v10

```

user@switch> show mld-snooping vlans vlan v10
VLAN          Interfaces Groups MRouters Receivers
v10              3      1      1      0      0

```

show mld-snooping vlans vlan vlan2 detail

```

user@switch> show mld-snooping vlans vlan vlan2 detail

VLAN: vlan2, Tag: 2, vlan-interface: vlan.2
  Interface: ge-0/0/2.0, untagged, Groups: 5
  Interface: ge-0/0/4.0, tagged, Groups: 3, Router

```

show mld-snooping vlans detail

```

user@switch> show mld-snooping vlans detail
VLAN: mld-vlan, Tag: 100
  Interface: ge-1/0/0.0, untagged, Groups: 0, Router
  Interface: ge-1/0/30.0, untagged, Groups: 1
  Interface: ge-1/0/33.0, untagged, Groups: 0
  Interface: ge-12/0/30.0, untagged, Groups: 0
VLAN: mld-vlan1, Tag: 200
  Interface: ge-1/0/31.0, untagged, Groups: 0
  Interface: ge-12/0/31.0, untagged, Groups: 1
  Interface: ae200.0, untagged, Groups: 0, Router

```


show multicast flow-map

| | |
|---|--|
| List of Syntax | Syntax on page 241 Syntax (EX Series Switch and the QFX Series) on page 241 |
| Syntax | show multicast flow-map <brief detail> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | show multicast flow-map <brief detail> |
| Release Information | Command introduced in Junos OS Release 8.2. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Display configuration information about IP multicast flow maps. |
| Options | none —Display configuration information about IP multicast flow maps on all systems. brief detail —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show multicast flow-map on page 242 show multicast flow-map detail on page 242 |
| Output Fields | Table 22 on page 241 describes the output fields for the show multicast flow-map command. Output fields are listed in the approximate order in which they appear. |

Table 22: show multicast flow-map Output Fields

| Field Name | Field Description | Levels of Output |
|----------------------|---|------------------|
| Name | Name of the flow map. | All levels |
| Policy | Name of the policy associated with the flow map. | All levels |
| Cache-timeout | Cache timeout value assigned to the flow map. | All levels |
| Bandwidth | Bandwidth setting associated with the flow map. | All levels |
| Adaptive | Whether or not adaptive mode is enabled for the flow map. | none |
| Flow-map | Name of the flow map. | detail |

Table 22: show multicast flow-map Output Fields (*continued*)

| Field Name | Field Description | Levels of Output |
|---------------------------|---|------------------|
| Adaptive Bandwidth | Whether or not adaptive mode is enabled for the flow map. | detail |
| Redundant Sources | Redundant sources defined for the same destination group. | detail |

Sample Output

show multicast flow-map

```

user@host> show multicast flow-map
Instance: master
Name          Policy          Cache timeout    Bandwidth Adaptive
map2          policy2         never            2000000 no
map1          policy1         60 seconds      2000000 no

```

Sample Output

show multicast flow-map detail

```

user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
  Policy:          policy1
  Cache Timeout:   600 seconds
  Bandwidth:       2000000
  Adaptive Bandwidth: yes
  Redundant Sources: 11.11.11.11
  Redundant Sources: 11.11.11.12
  Redundant Sources: 11.11.11.13

```


show multicast interface

| | |
|---|---|
| List of Syntax | Syntax on page 243 Syntax (EX Series Switch and the QFX Series) on page 243 |
| Syntax | <pre>show multicast interface <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | show multicast interface |
| Release Information | <p>Command introduced in Junos OS Release 8.3.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display bandwidth information about IP multicast interfaces. |
| Options | <p>none—Display all interfaces that have multicast configured.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show multicast interface on page 244 |
| Output Fields | <p>Table 23 on page 243 describes the output fields for the show multicast interface command. Output fields are listed in the approximate order in which they appear.</p> |

Table 23: show multicast interface Output Fields

| Field Name | Field Description |
|---|--|
| Interface | Name of the multicast interface. |
| Maximum bandwidth (bps) | Maximum bandwidth setting, in bits per second, for this interface. |
| Remaining bandwidth (bps) | Amount of bandwidth, in bits per second, remaining on the interface. |
| Mapped bandwidth deduction (bps) | <p>Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface.</p> <p>NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.</p> <p>This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |

Table 23: show multicast interface Output Fields (*continued*)

| Field Name | Field Description |
|--|--|
| Local bandwidth deduction (bps) | <p>Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface.</p> <p>NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.</p> <p>This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |
| Reverse OIF mapping | <p>State of the reverse OIF mapping feature (on or off).</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |
| Reverse OIF mapping no QoS adjustment | <p>State of the no QoS adjustment feature (on or off) for interfaces that are using reverse OIF mapping.</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |
| Leave timer | <p>Amount of time a mapped interface remains active after the last mapping ends.</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |
| No QoS adjustment | <p>State (on) of the no QoS adjustment feature when this feature is enabled.</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |

Sample Output

show multicast interface

```

user@host> show multicast interface
Interface           Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3             10000000                0
fe-0/0/3.210         10000000               -2000000
fe-0/0/3.220         100000000              100000000
fe-0/0/3.230         20000000               18000000
fe-0/0/2.200         100000000              100000000

```

show multicast minfo

| | |
|---------------------------------|--|
| Syntax | <code>show multicast minfo</code> <code><host></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Display configuration information about IP multicast networks, including neighboring multicast router addresses. |
| Options | none —Display configuration information about all multicast networks. host —(Optional) Display configuration information about a particular host. Replace <i>host</i> with a hostname or IP address. |
| Required Privilege Level | view |
| List of Sample Output | show multicast minfo on page 246 |
| Output Fields | Table 24 on page 245 describes the output fields for the show multicast minfo command. Output fields are listed in the approximate order in which they appear. |

Table 24: show multicast minfo Output Fields

| Field Name | Field Description |
|--------------------------------------|--|
| <i>source-address</i> | Query address, hostname (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor. |
| <i>ip-address-1—>ip-address-2</i> | Queried router interface address and directly attached neighbor interface address, respectively. |
| <i>(name or ip-address)</i> | Name or IP address of neighbor. |
| <i>[metric/threshold/type/flags]</i> | Neighbor's multicast profile: <ul style="list-style-type: none"> metric—Always has a value of 1, because minfo queries the directly connected interfaces of a device. threshold—Multicast threshold time-to-live (TTL). The range of values is 0 through 255. type—Multicast connection type: pim or tunnel. flags—Flags for this route: <ul style="list-style-type: none"> querier—Queried router is the designated router for the neighboring session. leaf—Link is a leaf in the multicast network. down—Link status indicator. |

Sample Output

show multicast mrinfo

```
user@host> show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
  192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
  0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```

show multicast next-hops

| | |
|---|---|
| List of Syntax | Syntax on page 247 Syntax (EX Series Switch and the QFX Series) on page 247 |
| Syntax | <pre>show multicast next-hops <brief detail> <identifier-number> <inet inet6> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show multicast next-hops <brief detail> <identifier-number> <inet inet6></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>detail option display of next-hop ID number introduced in Junos OS Release 11.1 for M Series and T Series routers and EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display the entries in the IP multicast next-hop table. |
| Options | <p>none—Display standard information about all entries in the multicast next-hop table for all supported address families.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>When you include the detail option on M Series and T Series routers and EX Series switches, the downstream interface name includes the next-hop ID number in parentheses, in the form fe-0/1/2.0-(1048574) where 1048574 is the next-hop ID number.</p> <p>identifier-number—(Optional) Show a particular next hop by ID number. The range of values is 1 through 65,535.</p> <p>inet inet6—(Optional) Display entries for IPv4 or IPv6 family addresses, respectively.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show multicast next-hops on page 248 show multicast next-hops (Bidirectional PIM on page 248 show multicast next-hops brief on page 249 show multicast next-hops detail on page 249 |

Output Fields Table 25 on page 248 describes the output fields for the **show multicast next-hops** command. Output fields are listed in the approximate order in which they appear.

Table 25: show multicast next-hops Output Fields

| Field Name | Field Description |
|--------------------------------|--|
| Family | Protocol family (such as INET). |
| ID | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine. |
| Refcount | Number of cache entries that are using this next hop. |
| KRefcount | Kernel reference count for the next hop. |
| Downstream interface | Interface names associated with each multicast next-hop ID. |
| Incoming interface list | List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM. |

Sample Output

show multicast next-hops

```
user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefcount Downstream interface
262142      4          2 so-1/0/0.0
262143      2          1 mt-1/1/0.49152
262148      2          1 mt-1/1/0.32769
```

show multicast next-hops (Bidirectional PIM)

```
user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefcount Downstream interface
2097151      8          4 ge-0/0/1.0

Family: INET6
ID      Refcount  KRefcount Downstream interface
2097157      2          1 ge-0/0/1.0

Family: Incoming interface list
ID      Refcount  KRefcount Downstream interface
513      5          2 lo0.0
                    ge-0/0/1.0
514      5          2 lo0.0
                    ge-0/0/1.0
                    xe-4/1/0.0
515      3          1 lo0.0
                    ge-0/0/1.0
                    xe-4/1/0.0
544      1          0 lo0.0
                    xe-4/1/0.0
```

show multicast next-hops brief

The output for the **show multicast next-hops brief** command is identical to that for the **show multicast next-hops** command. For sample output, see [show multicast next-hops on page 248](#).

show multicast next-hops detail

```
user@host> show multicast next-hops detail
Family: INET
ID          Refcount KRefCount Downstream interface
1048577      2          1 fe-0/1/2.0-(1048574)
              ge-0/2/3.0-(1048576)
```

show multicast pim-to-igmp-proxy

| | |
|---|---|
| List of Syntax | Syntax on page 250 Syntax (EX Series Switch and the QFX Series) on page 250 |
| Syntax | <pre>show multicast pim-to-igmp-proxy <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show multicast pim-to-igmp-proxy <instance <i>instance-name</i>></pre> |
| Release Information | <p>Command introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>instance option introduced in Junos OS Release 10.3.</p> <p>instance option introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy. |
| Options | <p>none—Display configuration information about PIM-to-IGMP message translation for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> Configuring PIM-to-IGMP and PIM-to-MLD Message Translation |
| List of Sample Output | show multicast pim-to-igmp-proxy on page 251 show multicast pim-to-igmp-proxy instance on page 251 |
| Output Fields | <p>Table 26 on page 250 describes the output fields for the show multicast pim-to-igmp-proxy command. Output fields are listed in the order in which they appear.</p> |

Table 26: show multicast pim-to-igmp-proxy Output Fields

| Field Name | Field Description |
|-----------------|---|
| Instance | Routing instance. Default instance is master (inet.0 routing table). |

Table 26: show multicast pim-to-igmp-proxy Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------|---|
| Proxy state | State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: enabled or disabled . |
| <i>interface-name</i> | Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured. |

Sample Output

show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

show multicast pim-to-igmp-proxy instance

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```

show multicast pim-to-mld-proxy

| | |
|---|---|
| List of Syntax | Syntax on page 252 Syntax (EX Series Switch and the QFX Series) on page 252 |
| Syntax | <pre>show multicast pim-to-mld-proxy <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show multicast pim-to-mld-proxy <instance <i>instance-name</i>></pre> |
| Release Information | <p>Command introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>instance option introduced in Junos OS Release 10.3.</p> <p>instance option introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy. |
| Options | <p>none—Display configuration information about PIM-to-MLD message translation for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show multicast pim-to-mld-proxy on page 253 show multicast pim-to-mld-proxy instance on page 253 |
| Output Fields | Table 27 on page 252 describes the output fields for the show multicast pim-to-mld-proxy command. Output fields are listed in the order in which they appear. |

Table 27: show multicast pim-to-mld-proxy Output Fields

| Field Name | Field Description |
|------------------------------|---|
| Proxy state | State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: enabled or disabled . |
| <i>interface-name</i> | Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured. |

Sample Output

show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

show multicast pim-to-mld-proxy instance

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```

show multicast route

List of Syntax [Syntax on page 254](#)
[Syntax \(EX Series Switch and the QFX Series\) on page 254](#)

Syntax `show multicast route`
 `<brief | detail | extensive | summary>`
 `<active | all | inactive>`
 `<group group>`
 `<inet | inet6>`
 `<instance instance name>`
 `<logical-system (all | logical-system-name)>`
 `<regular-expression>`
 `<source-prefix source-prefix>`

Syntax (EX Series Switch and the QFX Series) `show multicast route`
 `<brief | detail | extensive | summary>`
 `<active | all | inactive>`
 `<group group>`
 `<inet | inet6>`
 `<instance instance name>`
 `<regular-expression>`
 `<source-prefix source-prefix>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
 Command introduced in Junos OS Release 11.3 for the QFX Series.
 Support for bidirectional PIM added in Junos OS Release 12.1.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display the entries in the IP multicast forwarding table. You can display similar information with the **show route table inet.1** command.

Options **none**—Display standard information about all entries in the multicast forwarding table for all routing instances.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

active | all | inactive—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.

group group—(Optional) Display the cache entries for a particular group.

inet | inet6—(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.

instance instance-name—(Optional) Display entries in the multicast forwarding table for a specific multicast instance.

logical-system (all | logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.

regular-expression—(Optional) Display information about the multicast forwarding table entries that match a UNIX OS-style regular expression.

source-prefix source-prefix—(Optional) Display the cache entries for a particular source prefix.

Required Privilege Level view

Related Documentation

- *Example: Configuring Multicast-Only Fast Reroute in a PIM Domain*

List of Sample Output

- [show multicast route on page 256](#)
- [show multicast route \(Bidirectional PIM\) on page 257](#)
- [show multicast route brief on page 257](#)
- [show multicast route detail on page 258](#)
- [show multicast route extensive \(Bidirectional PIM\) on page 258](#)
- [show multicast route extensive \(Multicast-Only Fast Reroute\) on page 259](#)
- [show multicast route instance <instance-name> on page 259](#)
- [show multicast route summary on page 260](#)

Output Fields [Table 28 on page 255](#) describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

Table 28: show multicast route Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------------|---|------------------|
| family | IPv4 address family (INET) or IPv6 address family (INET6). | All levels |
| Group | Group address. For any-source multicast routes, for example for bidirectional PIM, the group address includes the prefix length. | All levels |
| Source | Prefix and length of the source as it is in the multicast forwarding table. | All levels |
| Incoming interface list | List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM. | All levels |
| Upstream interface | Name of the interface on which the packet with this source prefix is expected to arrive. | All levels |
| Upstream rpf interface list | When multicast-only fast reroute (MoFRR) is enabled, a PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. | All levels |
| Downstream interface list | List of interface names to which the packet with this source prefix is forwarded. | All levels |
| Number of outgoing interfaces | Total number of outgoing interfaces for each (S,G) entry. | extensive |

Table 28: show multicast route Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|---|-------------------|
| Session description | Name of the multicast session. | detail extensive |
| Statistics | Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobits per second packet forwarding statistic queries fails or times out, the statistics field displays Forwarding statistics are not available . NOTE: On QFX Series switches and OCX Series switches, this field does not report valid statistics. | detail extensive |
| Next-hop ID | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the show multicast nexthops command. | detail extensive |
| Incoming interface list ID | For bidirectional PIM, incoming interface list identifier. Identifiers for interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM. | detail extensive |
| Upstream protocol | The protocol that maintains the active multicast forwarding route for this group or source. When the show multicast route extensive command is used with the display-origin-protocol option, the field name is only Protocol and not Upstream Protocol . However, this field also displays the protocol that installed the active route. | detail extensive |
| Route type | Type of multicast route. Values can be (S,G) or (*G). | summary |
| Route state | Whether the group is Active or Inactive . | summary extensive |
| Route count | Number of multicast routes. | summary |
| Forwarding state | Whether the prefix is pruned or forwarding. | extensive |
| Cache lifetime/timeout | Number of seconds until the prefix is removed from the multicast forwarding table. A value of never indicates a permanent forwarding entry. A value of forever indicates routes that do not have keepalive times. | extensive |
| Wrong incoming interface notifications | Number of times that the upstream interface was not available. | extensive |
| Uptime | Time since the creation of a multicast route. | extensive |

Sample Output

show multicast route

```
user@host> show multicast route
```

```

Family: INET

Group: 228.0.0.0
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0

Group: 239.1.1.1
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0

Group: 239.1.1.1
  Source: 10.255.70.15/32
  Upstream interface: so-1/0/0.0
  Downstream interface list:
    mt-1/1/0.1081344

Family: INET6

```

show multicast route (Bidirectional PIM)

```

user@host> show multicast route
Family: INET

Group: 224.1.1.0/24
  Source: *
  Incoming interface list:
    lo0.0 ge-0/0/1.0
  Downstream interface list:
    ge-0/0/1.0

Group: 224.1.3.0/24
  Source: *
  Incoming interface list:
    lo0.0 ge-0/0/1.0 xe-4/1/0.0
  Downstream interface list:
    ge-0/0/1.0

Group: 225.1.1.0/24
  Source: *
  Incoming interface list:
    lo0.0 ge-0/0/1.0
  Downstream interface list:
    ge-0/0/1.0

Group: 225.1.3.0/24
  Source: *
  Incoming interface list:
    lo0.0 ge-0/0/1.0 xe-4/1/0.0
  Downstream interface list:
    ge-0/0/1.0
Family: INET6

```

show multicast route brief

The output for the **show multicast route brief** command is identical to that for the **show multicast route** command. For sample output, see [show multicast route on page 256](#) or [show multicast route \(Bidirectional PIM\) on page 257](#).

show multicast route detail

```
user@host> show multicast route detail
Family: INET

Group: 228.0.0.0
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Unknown
  Statistics: 8 kbps, 100 pps, 45272 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Administratively Scoped
  Statistics: 0 kbps, 0 pps, 13404 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.70.15/32
  Upstream interface: so-1/0/0.0
  Downstream interface list:
    mt-1/1/0.1081344
  Session description: Administratively Scoped
  Statistics: 46 kbps, 1000 pps, 921077 packets

  Next-hop ID: 262143
  Upstream protocol: PIM

Family: INET6
```

show multicast route extensive (Bidirectional PIM)

```
user@host> show multicast route extensive
Family: INET

Group: 224.1.1.0/24
  Source: *
  Incoming interface list:
    lo0.0 ge-0/0/1.0
  Downstream interface list:
    ge-0/0/1.0
  Number of outgoing interfaces: 1
  Session description: NOB Cross media facilities
  Statistics: 0 kbps, 0 pps, 0 packets
  Next-hop ID: 2097153
  Incoming interface list ID: 585
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0

Group: 224.1.3.0/24
```



```

Source: *
Incoming interface list:
  lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
  ge-0/0/1.0
Number of outgoing interfaces: 1
Session description: NOB Cross media facilities
Statistics: 0 kbps, 0 pps, 0 packets
Next-hop ID: 2097153
Incoming interface list ID: 589
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0

```

Family: INET6

show multicast route extensive (Multicast-Only Fast Reroute)

```
user@host> show multicast route extensive
```

Instance: master Family: INET

```

Group: 225.1.1.1
Source: 10.0.0.1/32
Upstream rpf interface list:
  fe-1/2/13.0 (P) fe-1/2/14.0 (B)
Downstream interface list:
  fe-1/2/15.0
Session description: Unknown
Forwarding statistics are not available
RPF Next-hop ID: 836
Next-hop ID: 1048585
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: 171 seconds
Wrong incoming interface notifications: 0
Uptime: 00:03:09

```

show multicast route instance <instance-name>

```
user@host> show multicast route instance v1 extensive
```

Instance: v1 Family: INET

```

Group: 224.1.1.1
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3

```

```

Group: 224.1.1.2
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3

```

```
Group: 224.1.1.3
```

```
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  1t-0/3/0.42 1t-0/3/0.46 1t-0/3/0.43
Number of outgoing interfaces: 3
```

```
Instance: v1 Family: INET6
```

show multicast route summary

```
user@host>show multicast route summary
Instance: master Family: INET
```

| Route type | Route state | Route count |
|------------|-------------|-------------|
| (S,G) | Active | 2 |
| (S,G) | Inactive | 3 |

```
Instance: master Family: INET6
```

show multicast rpf

| | |
|---|---|
| List of Syntax | Syntax on page 261 Syntax (EX Series Switch and the QFX Series) on page 261 |
| Syntax | <pre>show multicast rpf <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <prefix> <summary></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show multicast rpf <inet inet6> <instance <i>instance-name</i>> <prefix> <summary></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display information about multicast reverse-path-forwarding (RPF) calculations. |
| Options | <p>none—Display RPF calculation information for all supported address families.</p> <p>inet inet6—(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about multicast RPF calculations for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix—(Optional) Display the RPF calculation information for the specified prefix.</p> <p>summary—(Optional) Display a summary of all multicast RPF information.</p> |
| Required Privilege Level | view |
| List of Sample Output | show multicast rpf on page 262 show multicast rpf inet6 on page 263 show multicast rpf prefix on page 264 show multicast rpf summary on page 264 |

Output Fields Table 29 on page 262 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

Table 29: show multicast rpf Output Fields

| Field Name | Field Description |
|---------------|---|
| Instance | Name of the routing instance. (Displayed when multicast is configured within a routing instance.) |
| Source prefix | Prefix and length of the source as it exists in the multicast forwarding table. |
| Protocol | How the route was learned. |
| Interface | Upstream RPF interface. NOTE: The displayed interface information does not apply to bidirectional PIM RP addresses. This is because the show multicast rpf command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF interface information, always use the show pim join extensive command when bidirectional PIM is configured. |
| Neighbor | Upstream RPF neighbor. NOTE: The displayed neighbor information does not apply to bidirectional PIM. This is because the show multicast rpf command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF neighbor information, always use the show pim join extensive command when bidirectional PIM is configured. |

Sample Output

show multicast rpf

```

user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
  Protocol: Static

10.255.14.132/32
  Protocol: Direct
  Interface: lo0.0

10.255.245.91/32
  Protocol: IS-IS
  Interface: so-1/1/1.0
  Neighbor: 192.168.195.21

127.0.0.1/32
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0

```

```

Neighbor: 192.168.14.254

192.168.0.0/16
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

192.168.14.0/24
Protocol: Direct
Interface: fxp0.0

192.168.14.132/32
Protocol: Local

192.168.195.20/30
Protocol: Direct
Interface: so-1/1/1.0

192.168.195.22/32
Protocol: Local

192.168.195.36/30
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 192.168.195.21

```

show multicast rpf inet6

```

user@host> show multicast rpf inet6

Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128
  Protocol: Direct
  Interface: lo0.0

::10.255.245.91/128
  Protocol: IS-IS
  Interface: so-1/1/1.0
  Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.20/126
  Protocol: Direct
  Interface: so-1/1/1.0

::192.168.195.22/128
  Protocol: Local

::192.168.195.36/126
  Protocol: IS-IS
  Interface: so-1/1/1.0
  Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.76/126
  Protocol: Direct
  Interface: fe-2/2/0.0

::192.168.195.77/128
  Protocol: Local

```

```
fe80::/64
Protocol: Direct
Interface: so-1/1/1.0

fe80::290:69ff:fe0c:993a/128
Protocol: Local

fe80::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0

ff02::2/128
Protocol: PIM

ff02::d/128
Protocol: PIM
```

show multicast rpf prefix

```
user@host> show multicast rpf ff02::/16

Multicast RPF table: inet6.0, 13 entries

ff02::2/128
    Protocol: PIM

ff02::d/128
    Protocol: PIM

...
```

show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```

show multicast scope

| | |
|---|--|
| List of Syntax | Syntax on page 265 Syntax (EX Series Switch and the QFX Series) on page 265 |
| Syntax | <pre>show multicast scope <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show multicast scope <inet inet6> <instance <i>instance-name</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display administratively scoped IP multicast information. |
| Options | <p>none—Display standard information about administratively scoped multicast information for all supported address families in all routing instances.</p> <p>inet inet6—(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display administratively scoped information for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show multicast scope on page 266 show multicast scope inet on page 266 show multicast scope inet6 on page 266 |
| Output Fields | <p>Table 30 on page 265 describes the output fields for the show multicast scope command. Output fields are listed in the approximate order in which they appear.</p> |

Table 30: show multicast scope Output Fields

| Field Name | Field Description |
|--------------|--|
| Scope name | Name of the multicast scope. |
| Group Prefix | Range of multicast groups that are scoped. |

Table 30: show multicast scope Output Fields (*continued*)

| Field Name | Field Description |
|------------------------|---|
| Interface | Interface that is the boundary of the administrative scope. |
| Resolve Rejects | Number of kernel resolve rejects. |

Sample Output

show multicast scope

```
user@host> show multicast scope
```

| Scope name | Group Prefix | Interface | Resolve Rejects |
|------------|----------------|------------|-----------------|
| 232-net | 232.232.0.0/16 | fe-0/0/0.1 | 0 |
| local | 239.255.0.0/16 | fe-0/0/0.1 | 0 |
| local | ff05::/16 | fe-0/0/0.1 | 0 |
| larry | ff05::1234/128 | fe-0/0/0.1 | 0 |

show multicast scope inet

```
user@host> show multicast scope inet
```

| Scope name | Group Prefix | Interface | Resolve Rejects |
|------------|----------------|------------|-----------------|
| 232-net | 232.232.0.0/16 | fe-0/0/0.1 | 0 |
| local | 239.255.0.0/16 | fe-0/0/0.1 | 0 |

show multicast scope inet6

```
user@host> show multicast scope inet6
```

| Scope name | Group Prefix | Interface | Resolve Rejects |
|------------|----------------|------------|-----------------|
| local | ff05::/16 | fe-0/0/0.1 | 0 |
| larry | ff05::1234/128 | fe-0/0/0.1 | 0 |

show multicast sessions

| | |
|---|---|
| List of Syntax | Syntax on page 267 Syntax (EX Series Switch and the QFX Series) on page 267 |
| Syntax | show multicast sessions <brief detail extensive> <logical-system (all <i>logical-system-name</i>)> < <i>regular-expression</i> > |
| Syntax (EX Series Switch and the QFX Series) | show multicast sessions <brief detail extensive> < <i>regular-expression</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Display information about announced IP multicast sessions. |
| Options | none —Display standard information about all multicast sessions for all routing instances. brief detail extensive —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>regular-expression</i> —(Optional) Display information about announced sessions that match a UNIX-style regular expression. |
| Required Privilege Level | view |
| List of Sample Output | show multicast sessions on page 268 show multicast sessions regular-expression detail on page 268 |
| Output Fields | Table 31 on page 267 describes the output fields for the show multicast sessions command. Output fields are listed in the approximate order in which they appear. |

Table 31: show multicast sessions Output Fields

| Field Name | Field Description |
|---------------------|---|
| <i>session-name</i> | Name of the known announced multicast sessions. |

Sample Output

show multicast sessions

```

user@host> show multicast sessions
1-Department of Biological Sciences, LSU
...
Monterey Bay - DockCam
Monterey Bay - JettyCam
Monterey Bay - StandCam
Monterey DockCam
Monterey DockCam / ROV cam
...
NASA TV (MPEG-1)
...
UO Broadcast - NASA Videos - 25 Years of Progress
UO Broadcast - NASA Videos - Journey through the Solar System
UO Broadcast - NASA Videos - Life in the Universe
UO Broadcast - NASA Videos - Nasa and the Airplane
UO Broadcasts OPB's Oregon Story
UO DOD News Clips
UO Medical Management of Biological Casualties (1)
UO Medical Management of Biological Casualties (2)
UO Medical Management of Biological Casualties (3)
...
376 active sessions.

```

show multicast sessions regular-expression detail

```

user@host> show multicast sessions "NASA TV" detail
SDP Version: 0  Originated by: -@128.223.83.33
Session: NASA TV (MPEG-1)
Description: NASA television in MPEG-1 format, provided by Private University.
Please contact the UO if you have problems with this feed.
Email: Your Name Here <multicast@lists.private.edu>
Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 224.2.231.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 128.223.91.191 live
Attribute: fmtp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106
Connection Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
Attribute: rtpmap:103 L16/22050/2

```

Attribute: rtpmap:104 L16/22050

1 matching sessions.

show multicast usage

| | |
|---|--|
| List of Syntax | Syntax on page 270 Syntax (EX Series Switch and the QFX Series) on page 270 |
| Syntax | <code>show multicast usage</code> <code><brief detail></code> <code><inet inet6></code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Syntax (EX Series Switch and the QFX Series) | <code>show multicast usage</code> <code><brief detail></code> <code><inet inet6></code> <code><instance <i>instance-name</i>></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups. |
| Options | none —Display multicast usage information for all supported address families for all routing instances. brief detail —(Optional) Display the specified level of output. inet inet6 —(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively. instance <i>instance-name</i> —(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show multicast usage on page 271 show multicast usage brief on page 271 show multicast usage instance on page 271 show multicast usage detail on page 272 |
| Output Fields | Table 32 on page 271 describes the output fields for the show multicast usage command. Output fields are listed in the approximate order in which they appear. |

Table 32: show multicast usage Output Fields

| Field Name | Field Description |
|-----------------|--|
| Instance | Name of the routing instance. (Displayed when multicast is configured within a routing instance.) |
| Group | Group address. |
| Sources | Number of sources. |
| Packets | Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays unavailable . |
| Bytes | Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays unavailable . |
| Prefix | IP address. |
| /len | Prefix length. |
| Groups | Number of multicast groups. |

Sample Output

show multicast usage

```

user@host> show multicast usage
Group          Sources  Packets      Bytes
228.0.0.0      1        52847      4439148
239.1.1.1      2        13450      1125530

Prefix         /len  Groups  Packets      Bytes
10.255.14.144  /32   2        66254      5561304
10.255.70.15   /32   1         43        3374...
```

show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see [show multicast usage on page 271](#).

show multicast usage instance

```

user@host> show multicast usage instance VPN-A
Group          Sources  Packets      Bytes
224.2.127.254  1        5538      509496
224.0.1.39     1         13         624
224.0.1.40     1         13         624

Prefix         /len  Groups  Packets      Bytes
192.168.195.34 /32   1        5538      509496
10.255.14.30   /32   1         13         624
```

```
10.255.245.91 /32 1 13 624
...
```

show multicast usage detail

```
user@host> show multicast usage detail
```

| Group | Sources | Packets | Bytes |
|---|---------|---------|---------|
| 228.0.0.0 | 1 | 53159 | 4465356 |
| Source: 10.255.14.144 /32 Packets: 53159 Bytes: 4465356 | | | |
| 239.1.1.1 | 2 | 13450 | 1125530 |
| Source: 10.255.14.144 /32 Packets: 13407 Bytes: 1122156 | | | |
| Source: 10.255.70.15 /32 Packets: 43 Bytes: 3374 | | | |

| Prefix | /len | Groups | Packets | Bytes |
|------------------|------|----------------|----------------|---------|
| 10.255.14.144 | /32 | 2 | 66566 | 5587512 |
| Group: 228.0.0.0 | | Packets: 53159 | Bytes: 4465356 | |
| Group: 239.1.1.1 | | Packets: 13407 | Bytes: 1122156 | |
| 10.255.70.15 | /32 | 1 | 43 | 3374 |
| Group: 239.1.1.1 | | Packets: 43 | Bytes: 3374 | |

show pim bootstrap

| | |
|---|--|
| List of Syntax | Syntax on page 273 Syntax (EX Series Switch and the QFX Series) on page 273 |
| Syntax | <pre>show pim bootstrap <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show pim bootstrap <instance <i>instance-name</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>instance option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers. |
| Options | <p>none—Display PIM bootstrap router information for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show pim bootstrap on page 274 show pim bootstrap instance on page 274 |
| Output Fields | <p>Table 33 on page 273 describes the output fields for the show pim bootstrap command. Output fields are listed in the approximate order in which they appear.</p> |

Table 33: show pim bootstrap Output Fields

| Field Name | Field Description |
|----------------------|---|
| Instance | Name of the routing instance. |
| BSR | Bootstrap router. |
| Pri | Priority of the routing device as elected to be the bootstrap router. |
| Local address | Local routing device address. |

Table 33: show pim bootstrap Output Fields (*continued*)

| Field Name | Field Description |
|----------------|--|
| Pri | Local routing device address priority to be elected as the bootstrap router. |
| State | Local routing device election state: Candidate , Elected , or Ineligible . |
| Timeout | How long until the local routing device declares the bootstrap router to be unreachable, in seconds. |

Sample Output

show pim bootstrap

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

| BSR | Pri | Local address | Pri | State | Timeout |
|-------------------------|-----|-------------------------|-----|------------|---------|
| None | 0 | 10.255.71.46 | 0 | InEligible | 0 |
| feco:1:1:1:1:0:aff:785c | 34 | feco:1:1:1:1:0:aff:7c12 | 0 | InEligible | 0 |

show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A
```

| BSR | Pri | Local address | Pri | State | Timeout |
|------|-----|-----------------|-----|------------|---------|
| None | 0 | 192.168.196.105 | 0 | InEligible | 0 |

show pim interfaces

| | |
|---|---|
| List of Syntax | Syntax on page 275 Syntax (EX Series Switch and the QFX Series) on page 275 |
| Syntax | <pre>show pim interfaces <inet inet6> <instance (<i>instance-name</i> all)> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show pim interfaces <inet inet6> <instance (<i>instance-name</i> all)></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Support for the instance all option added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured. |
| Options | <p>none—Display interface information for all family addresses for the main instance.</p> <p>inet inet6—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance (<i>instance-name</i> all)—(Optional) Display information about interfaces for a specific PIM-enabled routing instance or for all routing instances.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show pim interfaces on page 276 |
| Output Fields | <p>Table 34 on page 275 describes the output fields for the show pim interfaces command. Output fields are listed in the approximate order in which they appear.</p> |

Table 34: show pim interfaces Output Fields

| Field Name | Field Description |
|-----------------|-------------------------------|
| Instance | Name of the routing instance. |
| Name | Interface name. |

Table 34: show pim interfaces Output Fields (*continued*)

| Field Name | Field Description |
|---------------------|---|
| State | State of the interface. The state also is displayed in the show interfaces command. |
| Mode | <p>PIM mode running on the interface:</p> <ul style="list-style-type: none"> • B—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers. • S—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic. • Dense—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.) • Sparse-Dense—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as dense is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as sparse is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules. <p>When sparse-dense mode is configured, the output includes both S and D. When bidirectional-sparse mode is configured, the output includes S and B. When bidirectional-sparse-dense mode is configured, the output includes B, S, and D.</p> |
| IP | Version number of the address family on the interface: 4 (IPv4) or 6 (IPv6). |
| V | PIM version running on the interface: 1 or 2. |
| State | <p>State of PIM on the interface:</p> <ul style="list-style-type: none"> • Active—Bidirectional mode is enabled on the interface and on all PIM neighbors. • DR—Designated router. • NotCap—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol. • NotDR—Not the designated router. • P2P—Point to point. |
| NbrCnt | Number of neighbors that have been seen on the interface. |
| JoinCnt(s,g) | Number of (s,g) join messages that have been seen on the interface. |
| JointCnt(*g) | Number of (*g) join messages that have been seen on the interface. |
| DR address | Address of the designated router. |

Sample Output

show pim interfaces

```
user@host> show pim interfaces
```

Stat = Status, V = Version, NbrCnt = Neighbor Count,
 S = Sparse, D = Dense, B = Bidirectional,
 DR = Designated Router, P2P = Point-to-point link,
 Active = Bidirectional is active, NotCap = Not Bidirectional Capable

| Name | Stat | Mode | IP | V | State | NbrCnt | JoinCnt(sg/*g) | DR address |
|----------------|------|------|----|---|--------------|--------|----------------|------------|
| ge-0/3/0.0 | Up | S | 4 | 2 | NotDR,NotCap | 1 | 0/0 | 40.0.0.3 |
| ge-0/3/3.50 | Up | S | 4 | 2 | DR,NotCap | 1 | 9901/100 | 50.0.0.2 |
| ge-0/3/3.51 | Up | S | 4 | 2 | DR,NotCap | 1 | 0/0 | 51.0.0.2 |
| pe-1/2/0.32769 | Up | S | 4 | 2 | P2P,NotCap | 0 | 0/0 | |

show pim join

List of Syntax [Syntax on page 278](#)
 [Syntax \(EX Series Switch and the QFX Series\) on page 278](#)

Syntax show pim join
 <brief | detail | extensive | summary>
 <bidirectional | dense | sparse>
 <exact>
 <inet | inet6>
 <instance *instance-name*>
 <logical-system (all | *logical-system-name*)>
 <range>
 <rp *ip-address/prefix* | source *ip-address/prefix*>
 <sg | star-g>

Syntax (EX Series Switch and the QFX Series) show pim join
 <brief | detail | extensive | summary>
 <dense | sparse>
 <exact>
 <inet | inet6>
 <instance *instance-name*>
 <range>
 <rp *ip-address/prefix* | source *ip-address/prefix*>
 <sg | star-g>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 summary option introduced in Junos OS Release 9.6.
 inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
 Support for bidirectional PIM added in Junos OS Release 12.1.
 Command introduced in Junos OS Release 11.3 for the QFX Series.
 Multiple new filter options introduced in Junos OS Release 13.2.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.

For bidirectional PIM, display information about PIM group ranges (*G-range) for each active bidirectional RP group range, in addition to each of the joined (*G) routes.

Options **none**—Display the standard information about PIM groups for all supported family addresses for all routing instances.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

bidirectional | dense | sparse—(Optional) Display information about PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.

exact—(Optional) Display information about only the group that exactly matches the specified group address.

inet | inet6—(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about groups for the specified PIM-enabled routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

range—(Optional) Address range of the group, specified as *prefix/prefix-length*.

rp *ip-address/prefix* | source *ip-address/prefix*—(Optional) Display information about the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

sg | star-g—(Optional) Display information about PIM (S,G) or (*,G) entries.

Required Privilege Level

view

Related Documentation

- [clear pim join on page 191](#)
- *Example: Configuring Multicast-Only Fast Reroute in a PIM Domain*
- *Example: Configuring Bidirectional PIM*
- *Example: Configuring PIM State Limits*

List of Sample Output

[show pim join summary on page 284](#)
[show pim join \(PIM Sparse Mode\) on page 284](#)
[show pim join \(Bidirectional PIM\) on page 284](#)
[show pim join inet6 on page 285](#)
[show pim join inet6 star-g on page 285](#)
[show pim join instance <instance-name> on page 285](#)
[show pim join detail on page 286](#)
[show pim join extensive \(PIM Sparse Mode\) on page 286](#)
[show pim join extensive \(Bidirectional PIM\) on page 287](#)
[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 288](#)
[show pim join instance <instance-name> extensive on page 289](#)
[show pim join extensive \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 289](#)
[show pim join extensive \(Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 290](#)
[show pim join summary on page 292](#)
[show pim join \(PIM Sparse Mode\) on page 292](#)
[show pim join \(Bidirectional PIM\) on page 293](#)
[show pim join inet6 on page 293](#)
[show pim join inet6 star-g on page 294](#)
[show pim join instance <instance-name> on page 294](#)
[show pim join detail on page 294](#)

[show pim join extensive \(PIM Sparse Mode\) on page 295](#)
[show pim join extensive \(Bidirectional PIM\) on page 296](#)
[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 297](#)
[show pim join instance <instance-name> extensive on page 297](#)
[show pim join extensive \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 298](#)
[show pim join extensive \(Multipoint LDP with Multicast-Only Fast Reroute\) on page 299](#)

Output Fields [Table 35 on page 280](#) describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

Table 35: show pim join Output Fields

| Field Name | Field Description | Level of Output |
|-----------------------------------|--|-------------------------------------|
| Instance | Name of the routing instance. | brief detail extensive summary none |
| Family | Name of the address family: inet (IPv4) or inet6 (IPv6). | brief detail extensive summary none |
| Route type | Type of multicast route: (S,G) or (*G). | summary |
| Route count | Number of (S,G) routes and number of (*G) routes. | summary |
| R | Rendezvous Point Tree. | brief detail extensive none |
| S | Sparse. | brief detail extensive none |
| W | Wildcard. | brief detail extensive none |
| Group | Group address. | brief detail extensive none |
| Bidirectional group prefix length | For bidirectional PIM, length of the IP prefix for RP group ranges. | All levels |
| Source | Multicast source: <ul style="list-style-type: none"> • * (wildcard value) • <i>ipv4-address</i> • <i>ipv6-address</i> | brief detail extensive none |
| RP | Rendezvous point for the PIM group. | brief detail extensive none |

Table 35: show pim join Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------------|---|------------------------------------|
| Flags | <p>PIM flags:</p> <ul style="list-style-type: none"> • bidirectional—Bidirectional mode entry. • dense—Dense mode entry. • rptree—Entry is on the rendezvous point tree. • sparse—Sparse mode entry. • spt—Entry is on the shortest-path tree for the source. • wildcard—Entry is on the shared tree. | brief detail extensive none |
| Upstream interface | <p>RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*,G).</p> <p>For bidirectional PIM, RP Link means that the interface is directly connected to a subnet that contains a phantom RP address.</p> <p>A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.</p> | brief detail extensive none |
| Upstream neighbor | <p>Information about the upstream neighbor: Direct, Local, Unknown, or a specific IP address.</p> <p>For bidirectional PIM, Direct means that the interface is directly connected to a subnet that contains a phantom RP address.</p> <p>The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.</p> | extensive |
| Active upstream interface | <p>When multicast-only fast reroute (MoFRR) is configured in a PIM domain, the upstream interface for the active path. A PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. Preference is given to two paths that do not converge to the same immediate upstream router. PIM installs appropriate multicast routes with upstream neighbors as RPF next hops with two (primary and backup) interfaces.</p> | extensive |
| Active upstream neighbor | <p>On the MoFRR primary path, the IP address of the neighbor that is directly connected to the active upstream interface.</p> | extensive |

Table 35: show pim join Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|--|------------------|
| MoFRR Backup upstream interface | <p>The MoFRR upstream interface that is used when the primary path fails.</p> <p>When the primary path fails, the backup path is upgraded to primary, and traffic is forwarded accordingly. If there are alternate paths available, a new backup path is calculated and the appropriate multicast route is updated or installed.</p> | extensive |
| Upstream state | <p>Information about the upstream interface:</p> <ul style="list-style-type: none"> • Join to RP—Sending a join to the rendezvous point. • Join to Source—Sending a join to the source. • Local RP—Sending neither join messages nor prune messages toward the RP, because this routing device is the rendezvous point. • Local Source—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device. • No Prune to RP—Automatically sent to RP when SPT and RPT are on the same path. • Prune to RP—Sending a prune to the rendezvous point. • Prune to Source—Sending a prune to the source. <p>NOTE: RP group range entries have None in the Upstream state field because RP group ranges do not trigger actual PIM join messages between routing devices.</p> | extensive |

Table 35: show pim join Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|------------------|
| Downstream neighbors | <p>Information about downstream interfaces:</p> <ul style="list-style-type: none"> • Interface—Interface name for the downstream neighbor. A pseudo PIM-SM interface appears for all IGMP-only interfaces. A pseudo multipoint LDP (M-LDP) interface appears on ingress root nodes in M-LDP point-to-multipoint LSPs with inband signaling. • Interface address—Address of the downstream neighbor. • State—Information about the downstream neighbor: join or prune. • Flags—PIM join flags: R (RPtree), S (Sparse), W (Wildcard), or zero. • Uptime—Time since the downstream interface joined the group. • Time since last Join—Time since the last join message was received from the downstream interface. • Time since last Prune—Time since the last prune message was received from the downstream interface. | extensive |
| Number of downstream interfaces | Total number of outgoing interfaces for each (S,G) entry. | extensive |
| Assert Timeout | Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null. | extensive |
| Keepalive timeout | Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, Keepalive timeout is Infinity . | extensive |
| Uptime | Time since the creation of (S,G) or (*,G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*,G) state. | extensive |
| Bidirectional accepting interfaces | <p>Interfaces on the routing device that forward bidirectional PIM traffic.</p> <p>The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (DF Winner), or the interface is the reverse path forwarding (RPF) interface toward the RP (RPF).</p> | extensive |

Sample Output

show pim join summary

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

Route type          Route count
(s,g)               2
(*,g)               1

Instance: PIM.master Family: INET6
```

show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
Bidirectional group prefix length: 24
Source: *
```

```

RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

show pim join inet6

```

user@host> show pim join inet6
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
Source: *
RP: ::46.0.0.13
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: ff04::e000:101
Source: ::1.1.1.1
Flags: sparse
Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
Source: ::1.1.1.1
Flags: sparse
Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
Source: ::1.1.1.2
Flags: sparse
Upstream interface: unknown (no neighbor)

```

show pim join inet6 star-g

```

user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
Source: *
RP: ::46.0.0.13
Flags: sparse,rptree,wildcard
Upstream interface: Local

```

show pim join instance <instance-name>

```

user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
Source: *
RP: 10.10.47.100

```

```

Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

show pim join detail

```

user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

show pim join extensive (PIM Sparse Mode)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 00:03:49
Downstream neighbors:
  Interface: so-1/0/0.0
    10.111.10.2 State: Join Flags: SRW Timeout: 174
    Uptime: 00:03:49 Time since last Join: 00:01:49
  Interface: mt-1/1/0.32768
    10.10.47.100 State: Join Flags: SRW Timeout: Infinity

```

```

        Uptime: 00:03:49 Time since last Join: 00:01:49
        Number of downstream interfaces: 2

Group: 239.1.1.1
  Source: 10.255.14.144
  Flags: sparse,spt
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local Source, Local RP
  Keepalive timeout: 344
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: so-1/0/0.0
      10.111.10.2 State: Join Flags: S Timeout: 174
      Uptime: 00:03:49 Time since last Prune: 00:01:49
    Interface: mt-1/1/0.32768
      10.10.47.100 State: Join Flags: S Timeout: Infinity
      Uptime: 00:03:49 Time since last Prune: 00:01:49
  Number of downstream interfaces: 2

Group: 239.1.1.1
  Source: 10.255.70.15
  Flags: sparse,spt
  Upstream interface: so-1/0/0.0
  Upstream neighbor: 10.111.10.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout: 344
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: Pseudo-GMP
      fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
    Interface: so-1/0/0.0 (pruned)
      10.111.10.2 State: Prune Flags: SR Timeout: 174
      Uptime: 00:03:49 Time since last Prune: 00:01:49
    Interface: mt-1/1/0.32768
      10.10.47.100 State: Join Flags: S Timeout: Infinity
      Uptime: 00:03:49 Time since last Prune: 00:01:49
  Number of downstream interfaces: 3

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

show pim join extensive (Bidirectional PIM)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0 (RPF)
    Interface: lo0.0 (DF Winner)
  Number of downstream interfaces: 0

```

```

Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
  Downstream neighbors:
    Interface: lt-1/0/10.24
      10.0.24.4 State: Join   RW   Timeout: 185
    Interface: lt-1/0/10.23
      10.0.23.3 State: Join   RW   Timeout: 184
  Number of downstream interfaces: 2

```

```

Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
    Interface: xe-4/1/0.0      (DF Winner)
  Number of downstream interfaces: 0

```

Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

```

Group: 224.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
    Interface: xe-4/1/0.0      (DF Winner)
  Number of downstream interfaces: 0

```

show pim join instance <instance-name> extensive

```

user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
Source: *
RP: 10.10.47.100
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 00:03:49
Downstream neighbors:
  Interface: mt-1/1/0.32768
    10.10.47.101 State: Join Flags: SRW Timeout: 156
    Uptime: 00:03:49 Time since last Join: 00:01:49
  Number of downstream interfaces: 1

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0
Upstream neighbor: 10.111.30.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0
Upstream neighbor: 10.111.20.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52

```

show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:55
Downstream neighbors:
  Interface: Pseudo-MLDP
  Interface: lt-1/2/0.25
    1.2.5.2 State: Join Flags: S Timeout: Infinity
    Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
Source: 192.168.219.11
Flags: sparse,spt

```

```
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:41
Downstream neighbors:
    Interface: Pseudo-MLDP

Group: 232.1.1.3
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:41
Downstream neighbors:
    Interface: Pseudo-MLDP

Group: 232.2.2.2
Source: 1.2.7.7
Flags: sparse,spt
Upstream interface: lt-1/2/0.27
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:25
Downstream neighbors:
    Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
Source: abcd::1:2:7:7
Flags: sparse,spt
Upstream interface: lt-1/2/0.27
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:26
Downstream neighbors:
    Interface: Pseudo-MLDP
```

show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 227.1.1.1
Source: *
RP: 1.1.1.1
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 11:31:33
Downstream neighbors:
    Interface: fe-1/3/0.0
    192.168.209.9 State: Join Flags: SRW Timeout: Infinity
```


Uptime: 11:31:33 Time since last Join: 11:31:32

Group: 232.1.1.1

Source: 192.168.219.11
 Flags: sparse,spt
 Upstream protocol: MLDP
 Upstream interface: Pseudo MLDP
 Upstream neighbor: MLDP LSP root <1.1.1.2>
 Upstream state: Join to Source
 Keepalive timeout:
 Uptime: 11:31:32
 Downstream neighbors:
 Interface: so-0/1/3.0
 192.168.92.9 State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:30 Time since last Join: 11:31:30
 Downstream neighbors:
 Interface: fe-1/3/0.0
 192.168.209.9 State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.2

Source: 192.168.219.11
 Flags: sparse,spt
 Upstream protocol: MLDP
 Upstream interface: Pseudo MLDP
 Upstream neighbor: MLDP LSP root <1.1.1.2>
 Upstream state: Join to Source
 Keepalive timeout:
 Uptime: 11:31:32
 Downstream neighbors:
 Interface: so-0/1/3.0
 192.168.92.9 State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:30 Time since last Join: 11:31:30
 Downstream neighbors:
 Interface: lt-1/2/0.14
 1.1.4.4 State: Join Flags: S Timeout: 177
 Uptime: 11:30:33 Time since last Join: 00:00:33
 Downstream neighbors:
 Interface: fe-1/3/0.0
 192.168.209.9 State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.3

Source: 192.168.219.11
 Flags: sparse,spt
 Upstream protocol: MLDP
 Upstream interface: Pseudo MLDP
 Upstream neighbor: MLDP LSP root <1.1.1.2>
 Upstream state: Join to Source
 Keepalive timeout:
 Uptime: 11:31:32
 Downstream neighbors:
 Interface: fe-1/3/0.0
 192.168.209.9 State: Join Flags: S Timeout: Infinity
 Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.2.2.2

Source: 1.2.7.7
 Flags: sparse,spt
 Upstream protocol: MLDP
 Upstream interface: Pseudo MLDP

```
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 11:31:30
Downstream neighbors:
  Interface: so-0/1/3.0
    192.168.92.9 State: Join Flags: S   Timeout: Infinity
    Uptime: 11:31:30 Time since last Join: 11:31:30

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
Source: abcd::1:2:7:7
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 11:31:32
Downstream neighbors:
  Interface: fe-1/3/0.0
    fe80::21f:12ff:fea5:c4db State: Join Flags: S   Timeout: Infinity
    Uptime: 11:31:32 Time since last Join: 11:31:32
```

Sample Output

show pim join summary

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

Route type          Route count
(s,g)                2
(*,g)                1

Instance: PIM.master Family: INET6
```

show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0
```

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join inet6

```
user@host> show pim join inet6
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
  Source: *
  RP: ::46.0.0.13
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

Group: ff04::e000:101
  Source: ::1.1.1.1
  Flags: sparse
  Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
  Source: ::1.1.1.1
  Flags: sparse
  Upstream interface: unknown (no neighbor)
```

```
Group: ff04::e800:101
Source: ::1.1.1.2
Flags: sparse
Upstream interface: unknown (no neighbor)
```

show pim join inet6 star-g

```
user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
Source: *
RP: ::46.0.0.13
Flags: sparse,rptree,wildcard
Upstream interface: Local
```

show pim join instance <instance-name>

```
user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
Source: *
RP: 10.10.47.100
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join detail

```
user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
```

```
Flags: sparse,spt
Upstream interface: so-1/0/0.0
```

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: *
  RP: 10.255.14.144
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local RP
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: so-1/0/0.0
      10.111.10.2 State: Join Flags: SRW Timeout: 174
      Uptime: 00:03:49 Time since last Join: 00:01:49
    Interface: mt-1/1/0.32768
      10.10.47.100 State: Join Flags: SRW Timeout: Infinity
      Uptime: 00:03:49 Time since last Join: 00:01:49
  Number of downstream interfaces: 2

Group: 239.1.1.1
  Source: 10.255.14.144
  Flags: sparse,spt
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local Source, Local RP
  Keepalive timeout: 344
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: so-1/0/0.0
      10.111.10.2 State: Join Flags: S Timeout: 174
      Uptime: 00:03:49 Time since last Prune: 00:01:49
    Interface: mt-1/1/0.32768
      10.10.47.100 State: Join Flags: S Timeout: Infinity
      Uptime: 00:03:49 Time since last Prune: 00:01:49
  Number of downstream interfaces: 2

Group: 239.1.1.1
  Source: 10.255.70.15
  Flags: sparse,spt
  Upstream interface: so-1/0/0.0
  Upstream neighbor: 10.111.10.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout: 344
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: Pseudo-GMP
      fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
    Interface: so-1/0/0.0 (pruned)
      10.111.10.2 State: Prune Flags: SR Timeout: 174
      Uptime: 00:03:49 Time since last Prune: 00:01:49
    Interface: mt-1/1/0.32768
```

```

10.10.47.100 State: Join Flags: S   Timeout: Infinity
Uptime: 00:03:49 Time since last Prune: 00:01:49
Number of downstream interfaces: 3

```

```

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

show pim join extensive (Bidirectional PIM)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

```

Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
  Number of downstream interfaces: 0

Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
  Downstream neighbors:
    Interface: lt-1/0/10.24
      10.0.24.4 State: Join   RW   Timeout: 185
    Interface: lt-1/0/10.23
      10.0.23.3 State: Join   RW   Timeout: 184
  Number of downstream interfaces: 2

Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
    Interface: xe-4/1/0.0      (DF Winner)
  Number of downstream interfaces: 0

```

Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
    Interface: xe-4/1/0.0     (DF Winner)
  Number of downstream interfaces: 0
```

show pim join instance <instance-name> extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
  Source: *
  RP: 10.10.47.100
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local RP
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: mt-1/1/0.32768
      10.10.47.101 State: Join Flags: SRW Timeout: 156
      Uptime: 00:03:49 Time since last Join: 00:01:49
  Number of downstream interfaces: 1

Group: 235.1.1.2
  Source: 192.168.195.74
  Flags: sparse,spt
  Upstream interface: at-0/3/1.0
  Upstream neighbor: 10.111.30.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout: 156
  Uptime: 00:14:52

Group: 235.1.1.2
  Source: 192.168.195.169
  Flags: sparse
  Upstream interface: so-1/0/1.0
  Upstream neighbor: 10.111.20.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout: 156
  Uptime: 00:14:52
```

show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream interface: fe-1/3/1.0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 11:27:55
  Downstream neighbors:
    Interface: Pseudo-MLDP
    Interface: lt-1/2/0.25
      1.2.5.2 State: Join Flags: S   Timeout: Infinity
      Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream interface: fe-1/3/1.0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 11:27:41
  Downstream neighbors:
    Interface: Pseudo-MLDP

Group: 232.1.1.3
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream interface: fe-1/3/1.0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 11:27:41
  Downstream neighbors:
    Interface: Pseudo-MLDP

Group: 232.2.2.2
  Source: 1.2.7.7
  Flags: sparse,spt
  Upstream interface: lt-1/2/0.27
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 11:27:25
  Downstream neighbors:
    Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
  Source: abcd::1:2:7:7
  Flags: sparse,spt
  Upstream interface: lt-1/2/0.27
  Upstream neighbor: Direct
```



```

Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:26
Downstream neighbors:
    Interface: Pseudo-MLDP

```

show pim join extensive (Multipoint LDP with Multicast-Only Fast Reroute)

```

user@host> show pim join 225.1.1.1 extensive sg
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.1.1.1
  Source: 10.0.0.1
  Flags: sparse,spt
  Active upstream interface: fe-1/2/13.0
  Active upstream neighbor: 10.0.0.9
  MoFRR Backup upstream interface: fe-1/2/14.0
  MoFRR Backup upstream neighbor: 10.0.0.21
  Upstream state: Join to Source, No Prune to RP
  Keepalive timeout: 354
  Uptime: 00:00:06
  Downstream neighbors:
    Interface: fe-1/2/15.0
      10.0.0.13 State: Join Flags: S Timeout: Infinity
      Uptime: 00:00:06 Time since last Join: 00:00:06
  Number of downstream interfaces: 1

```

show pim neighbors

| | |
|---|--|
| List of Syntax | Syntax on page 300 Syntax (EX Series Switch and the QFX Series) on page 300 |
| Syntax | <pre>show pim neighbors <brief detail> <inet inet6> <instance (<i>instance-name</i> all)> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show pim neighbors <brief detail> <inet inet6> <instance (<i>instance-name</i> all)></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Support for the instance all option added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display information about Protocol Independent Multicast (PIM) neighbors. |
| Options | <p>none—(Same as brief) Display standard information about PIM neighbors for all supported family addresses for the main instance.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>inet inet6—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.</p> <p>instance (<i>instance-name</i> all)—(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show pim neighbors on page 302 show pim neighbors brief on page 302 show pim neighbors instance on page 302 show pim neighbors detail on page 302 show pim neighbors detail (With BFD) on page 303 |
| Output Fields | Table 36 on page 301 describes the output fields for the show pim neighbors command. Output fields are listed in the approximate order in which they appear. |

Table 36: show pim neighbors Output Fields

| Field Name | Field Description | Level of Output |
|--|--|-------------------|
| Instance | Name of the routing instance. | All levels |
| Interface | Interface through which the neighbor is reachable. | All levels |
| Neighbor addr | Address of the neighboring PIM routing device. | All levels |
| IP | IP version: 4 or 6. | All levels |
| V | PIM version running on the neighbor: 1 or 2. | All levels |
| Mode | PIM mode of the neighbor: Sparse , Dense , SparseDense , or Unknown . When the neighbor is running PIM version 2, this mode is always Unknown . | All levels |
| Option | Can be one or more of the following: <ul style="list-style-type: none"> • B—Bidirectional Capable. • G—Generation Identifier. • H—Hello Option Holdtime. • L—Hello Option LAN Prune Delay. • P—Hello Option DR Priority. • T—Tracking bit. | brief none |
| Uptime | Time the neighbor has been operational since the PIM process was last initialized, in the format dd:hh:mm:ss ago for less than a week and nwnd:hh:mm:ss ago for more than a week. | All levels |
| Address | Address of the neighboring PIM routing device. | detail |
| BFD | Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: Enabled , Operational state is up , or Disabled . | detail |
| Hello Option Holdtime | Time for which the neighbor is available, in seconds. The range of values is 0 through 65,535. | detail |
| Hello Default Holdtime | Default holdtime and the time remaining if the holdtime option is not in the received hello message. | detail |
| Hello Option DR Priority | Designated router election priority. The range of values is 0 through 255. | detail |
| Hello Option Generation ID | 9-digit or 10-digit number used to tag hello messages. | detail |
| Hello Option Bi-Directional PIM supported | Neighbor can process bidirectional PIM messages. | detail |
| Hello Option LAN Prune Delay | Time to wait before the neighbor receives prune messages, in the format delay nnn ms override nnnn ms . | detail |

Table 36: show pim neighbors Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------|---|-----------------|
| Join Suppression supported | Neighbor is capable of join suppression. | detail |
| Rx Join | Information about joins received from the neighbor. <ul style="list-style-type: none"> Group—Group addresses in the join message. Source—Address of the source in the join message. Timeout—Time for which the join is valid. | detail |

Sample Output

show pim neighbors

```

user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking bit

Interface      IP V Mode      Option      Uptime Neighbor addr
so-1/0/0.0      4 2            HPLG        00:07:10 10.111.10.2

```

show pim neighbors brief

The output for the **show pim neighbors brief** command is identical to that for the **show pim neighbors** command. For sample output, see [show pim neighbors on page 302](#).

show pim neighbors instance

```

user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking bit

Interface      IP V Mode      Option      Uptime Neighbor addr
at-0/3/1.0      4 2            HPLG        00:07:54 10.111.30.2
mt-1/1/0.32768  4 2            HPLG        00:07:22 10.10.47.101
so-1/0/1.0      4 2            HPLG        00:07:50 10.111.20.2

```

show pim neighbors detail

```

user@host> show pim neighbors detail
Instance: PIM.master
Interface: ge-0/0/1.0

Address: 10.10.1.1, IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, ts
Join Count: 2
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 2053759302
Hello Option Bi-Directional PIM supported
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

```

```

Address: 10.10.1.2, IPv4, PIM v2, sg Join Count: 0, tsg Join Count: 2
  BFD: Disabled
  Hello Option Holdtime: 105 seconds 93 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1734018161
  Hello Option Bi-Directional PIM supported
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Join Suppression supported

```

```
Interface: lo0.0
```

```

Address: 10.255.179.246, IPv4, PIM v2, Mode: SparseDense, sg Join Count:
0, tsg Join Count: 0
  Hello Option Holdtime: 65535 seconds
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1997462267
  Hello Option Bi-Directional PIM supported
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Join Suppression supported

```

show pim neighbors detail (With BFD)

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

```
Instance: PIM.master
```

```
Interface: fe-1/0/0.0
```

```

Address: 192.168.11.1, IPv4, PIM v2, Mode: Sparse
  Hello Option Holdtime: 65535 seconds
  Hello Option DR Priority: 1
  Hello Option Generation ID: 836607909
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

```

```

Address: 192.168.11.2, IPv4, PIM v2
  BFD: Enabled, Operational state is up
  Hello Default Holdtime: 105 seconds 104 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1907549685
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

```

```
Interface: fe-1/0/1.0
```

```

Address: 192.168.12.1, IPv4, PIM v2
  BFD: Disabled
  Hello Default Holdtime: 105 seconds 80 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1971554705
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

```

show pim rps

| | |
|--|--|
| List of Syntax | Syntax on page 304 Syntax (EX Series Switch and the QFX Series) on page 304 |
| Syntax | <pre>show pim rps <brief detail extensive> <group-address> <inet inet6> <instance instance-name> <logical-system (all logical-system-name)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show pim rps <brief detail extensive> <group-address> <inet inet6> <instance instance-name></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs). |
| Options | <p>none—Display standard information about PIM RPs for all groups and family addresses for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>group-address—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.</p> <p>inet inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance instance-name—(Optional) Display information about RPs for a specific PIM-enabled routing instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring Bidirectional PIM |
| List of Sample Output | show pim rps on page 307 show pim rps brief on page 307 |

[show pim rps <group-address> on page 307](#)
[show pim rps <group-address> on page 307](#)
[show pim rps <group-address> \(Bidirectional PIM\) on page 308](#)
[show pim rps <group-address> \(PIM Dense Mode\) on page 308](#)
[show pim rps <group-address> \(SSM Range Without asm-override-ssm Configured\) on page 308](#)
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Sparse-Mode RP\) on page 308](#)
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Bidirectional RP\) on page 308](#)
[show pim rps instance on page 308](#)
[show pim rps extensive \(PIM Sparse Mode\) on page 308](#)
[show pim rps extensive \(Bidirectional PIM\) on page 309](#)
[show pim rps extensive \(PIM Anycast RP in Use\) on page 309](#)

Output Fields [Table 37 on page 305](#) describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

Table 37: show pim rps Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------------|--|-------------------------|
| Instance | Name of the routing instance. | All levels |
| Family or Address family | Name of the address family: inet (IPv4) or inet6 (IPv6). | All levels |
| RP address | Address of the rendezvous point. | All levels |
| Type | Type of RP: <ul style="list-style-type: none"> • auto-rp—Address of the RP known through the Auto-RP protocol. • bootstrap—Address of the RP known through the bootstrap router protocol (BSR). • embedded—Address of the RP known through an embedded RP (IPv6). • static—Address of RP known through static configuration. | brief none |
| Holdtime | How long to keep the RP active, with time remaining, in seconds. | All levels |
| Timeout | How long until the local routing device determines the RP to be unreachable, in seconds. | All levels |
| Groups | Number of groups currently using this RP. | All levels |
| Group prefixes | Addresses of groups that this RP can span. | brief none |
| Learned via | Address and method by which the RP was learned. | detail extensive |
| Mode | The PIM mode of the RP: bidirectional or sparse. If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats. | All levels |

Table 37: show pim rps Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------|--|--|
| Time Active | How long the RP has been active, in the format <i>hh:mm:ss</i> . | detail extensive |
| Device Index | Index value of the order in which Junos OS finds and initializes the interface. For bidirectional RPs, the Device Index output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces. | detail extensive |
| Subunit | Logical unit number of the interface. For bidirectional RPs, the Subunit output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces. | detail extensive |
| Interface | Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively. For bidirectional RPs, the Interface output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces. | detail extensive |
| Group Ranges | Addresses of groups that this RP spans. | detail extensive <i>group-address</i> |
| Active groups using RP | Number of groups currently using this RP. | detail extensive |
| total | Total number of active groups for this RP. | detail extensive |
| Register State for RP | Current register state for each group: <ul style="list-style-type: none"> • Group—Multicast group address. • Source—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively: • First Hop—PIM-designated routing device that sent the Register message (the source address in the IP header). • RP Address—RP to which the Register message was sent (the destination address in the IP header). • State: <ul style="list-style-type: none"> On the designated router: <ul style="list-style-type: none"> • Send—Sending Register messages. • Probe—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages. • Suppress—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to Probe state. On the RP: <ul style="list-style-type: none"> • Receive—Receiving Register messages. | extensive |
| Anycast-PIM rpset | If anycast RP is configured, the addresses of the RPs in the set. | extensive |

Table 37: show pim rps Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|----------------------|
| Anycast-PIM local address used | If anycast RP is configured, the local address used by the RP. | extensive |
| Anycast-PIM Register State | If anycast RP is configured, the current register state for each group: <ul style="list-style-type: none"> • Group—Multicast group address. • Source—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively. • Origin—How the information was obtained: <ul style="list-style-type: none"> • DIRECT—From a local attachment • MSDP—From the Multicast Source Discovery Protocol (MSDP) • DR—From the designated router | extensive |
| RP selected | For sparse mode and bidirectional mode, the identity of the RP for the specified group address. | <i>group-address</i> |

Sample Output

show pim rps

```

user@host> show pim rps
Instance: PIM.master

Address-family INET
RP address      Type      Mode   Holdtime Timeout Groups  Group prefixes
100.100.100.100 auto-rp   sparse    150     146      0 235.0.0.0/8
                                     235.100.100.0/24
200.200.200.200 auto-rp   sparse    150     146      0 224.0.0.0/4

address-family INET6

```

show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see [show pim rps on page 307](#).

show pim rps <group-address>

```

user@host> show pim rps 235.100.100.0
Instance: PIM.master
Instance: PIM.master

RP selected: 100.100.100.100

```

show pim rps <group-address>

```

user@host> show pim rps 235.100.100.0
Instance: PIM.master
Instance: PIM.master

RP selected: 100.100.100.100

```

show pim rps <group-address> (Bidirectional PIM)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

224.1.0.0/16
    11.4.12.75 (Bidirectional)

RP selected: 11.4.12.75
```

show pim rps <group-address> (PIM Dense Mode)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Dense Mode active for group 224.1.1.1
```

show pim rps <group-address> (SSM Range Without asm-override-ssm Configured)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active for group 224.1.1.1
```

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16
    11.4.12.75

RP selected: 11.4.12.75
```

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16
    11.4.12.75 (Bidirectional)

RP selected: (null)
```

show pim rps instance

```
user@host> show pim rps instance VPN-A
Instance: PIM.VPN-A
Address family INET
RP address          Type      Holdtime Timeout Groups Group prefixes
10.10.47.100        static    0        None    1 224.0.0.0/4

Address family INET6
```

show pim rps extensive (PIM Sparse Mode)

```
user@host> show pim rps extensive
```

Instance: PIM.master

Family: INET
 RP: 10.255.245.91
 Learned via: static configuration
 Time Active: 00:05:48
 Holdtime: 45 with 36 remaining
 Device Index: 122
 Subunit: 32768
 Interface: pd-6/0/0.32768
 Group Ranges:
 224.0.0.0/4, 36s remaining
 Active groups using RP:
 225.1.1.1

total 1 groups active

Register State for RP:

| Group | Source | FirstHop | RP Address | State | Timeout |
|-----------|----------------|---------------|---------------|---------|---------|
| 225.1.1.1 | 192.168.195.78 | 10.255.14.132 | 10.255.245.91 | Receive | 0 |

show pim rps extensive (Bidirectional PIM)

user@host> show pim rps extensive

Instance: PIM.master
 Address family INET

RP: 10.10.1.3
 Learned via: static configuration
 Mode: Bidirectional
 Time Active: 01:58:07
 Holdtime: 150
 Group Ranges:
 224.1.3.0/24
 225.1.3.0/24

RP: 10.10.13.2
 Learned via: static configuration
 Mode: Bidirectional
 Time Active: 01:58:07
 Holdtime: 150
 Group Ranges:
 224.1.1.0/24
 225.1.1.0/24

show pim rps extensive (PIM Anycast RP in Use)

user@host> show pim rps extensive

Instance: PIM.master

Family: INET
 RP: 10.10.10.2
 Learned via: static configuration
 Time Active: 00:54:52
 Holdtime: 0
 Device Index: 130
 Subunit: 32769
 Interface: pimd.32769
 Group Ranges:
 224.0.0.0/4
 Active groups using RP:

224.10.10.10

total 1 groups active

Anycast-PIM rpset:

10.100.111.34

10.100.111.17

10.100.111.55

Anycast-PIM local address used: 10.100.111.1

Anycast-PIM Register State:

| Group | Source | Origin |
|--------------|------------|--------|
| 224.1.1.1 | 10.10.95.2 | DIRECT |
| 224.1.1.2 | 10.10.95.2 | DIRECT |
| 224.10.10.10 | 10.10.70.1 | MSDP |
| 224.10.10.11 | 10.10.70.1 | MSDP |
| 224.20.20.1 | 10.10.71.1 | DR |

Address family INET6

Anycast-PIM rpset:

ab::1

ab::2

Anycast-PIM local address used: cd::1

Anycast-PIM Register State:

| Group | Source | Origin |
|---------------|--------------|--------|
| ::224.1.1.1 | ::10.10.95.2 | DIRECT |
| ::224.1.1.2 | ::10.10.95.2 | DIRECT |
| ::224.20.20.1 | ::10.10.71.1 | DR |

show pim source

| | |
|---|--|
| List of Syntax | Syntax on page 311 Syntax (EX Series Switch and the QFX Series) on page 311 |
| Syntax | <pre>show pim source <brief detail> <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <source-prefix></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show pim source <brief detail> <inet inet6> <instance <i>instance-name</i>> <source-prefix></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state. |
| Options | <p>none—Display standard information about the PIM RPF state for all supported family addresses for all routing instances.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>inet inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>source-prefix—(Optional) Display the state for source RPF states in the given range.</p> |
| Required Privilege Level | view |
| List of Sample Output | show pim source on page 312 show pim source brief on page 312 show pim source detail on page 312 show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 313 |

Output Fields [Table 38 on page 312](#) describes the output fields for the **show pim source** command. Output fields are listed in the approximate order in which they appear.

Table 38: show pim source Output Fields

| Field Name | Field Description |
|---------------------------|---|
| Instance | Name of the routing instance. |
| Source | Address of the source or reverse path. |
| Prefix/length | Prefix and prefix length for the route used to reach the RPF address. |
| Upstream Protocol | Protocol toward the source address. |
| Upstream interface | RPF interface toward the source address. A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling. |
| Upstream Neighbor | Address of the RPF neighbor used to reach the source address. The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling. |

Sample Output

show pim source

```

user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
  Prefix 10.255.14.144/32
  Upstream interface Local
  Upstream neighbor Local

Source 10.255.70.15
  Prefix 10.255.70.15/32
  Upstream interface so-1/0/0.0
  Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6

```

show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see [show pim source on page 312](#).

show pim source detail

```

user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
  Prefix 10.255.14.144/32
  Upstream interface Local

```

```

Upstream neighbor Local
Active groups:228.0.0.0
239.1.1.1
239.1.1.1

Source 10.255.70.15
Prefix 10.255.70.15/32
Upstream interface so-1/0/0.0
Upstream neighbor 10.111.10.2
Active groups:239.1.1.1

Instance: PIM.master Family: INET6

```

show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show pim source
Instance: PIM.master Family: INET

Source 1.1.1.1
Prefix 1.1.1.1/32
Upstream interface Local
Upstream neighbor Local

Source 1.2.7.7
Prefix 1.2.7.0/24
Upstream protocol MLDP
Upstream interface Pseudo MLDP
Upstream neighbor MLDP LSP root <1.1.1.2>

Source 192.168.219.11
Prefix 192.168.219.0/28
Upstream protocol MLDP
Upstream interface Pseudo MLDP
Upstream neighbor MLDP LSP root <1.1.1.2>

Instance: PIM.master Family: INET6

Source abcd::1:2:7:7
Prefix abcd::1:2:7:0/120
Upstream protocol MLDP
Upstream interface Pseudo MLDP
Upstream neighbor MLDP LSP root <1.1.1.2>

```

show pim statistics

| | |
|---|---|
| List of Syntax | Syntax on page 314 Syntax (EX Series Switch and the QFX Series) on page 314 |
| Syntax | <pre>show pim statistics <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show pim statistics <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display Protocol Independent Multicast (PIM) statistics. |
| Options | <p>none—Display PIM statistics.</p> <p>inet inet6—(Optional) Display IPv4 or IPv6 PIM statistics, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM).</p> <p>interface <i>interface-name</i>—(Optional) Display statistics about the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• clear pim statistics on page 195 |
| List of Sample Output | <p>show pim statistics on page 321</p> <p>show pim statistics inet interface <interface-name> on page 323</p> <p>show pim statistics inet6 interface <interface-name> on page 323</p> <p>show pim statistics instance <instance-name> on page 324</p> <p>show pim statistics interface <interface-name> on page 325</p> |
| Output Fields | <p>Table 39 on page 315 describes the output fields for the show pim statistics command. Output fields are listed in the approximate order in which they appear.</p> |

Table 39: show pim statistics Output Fields

| Field Name | Field Description |
|------------------|---|
| Instance | <p>Name of the routing instance.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> • inet interface <i>interface-name</i> • inet6 interface <i>interface-name</i> • interface <i>interface-name</i> |
| Family | <p>Output is for IPv4 or IPv6 PIM statistics. INET indicates IPv4 statistics, and INET6 indicates IPv6 statistics.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> • inet interface <i>interface-name</i> • inet6 interface <i>interface-name</i> • interface <i>interface-name</i> |
| PIM statistics | PIM statistics for all interfaces or for the specified interface. |
| PIM message type | Message type for which statistics are displayed. |
| Received | Number of received statistics. |
| Sent | Number of messages sent of a certain type. |
| Rx errors | Number of received packets that contained errors. |
| V2 Hello | PIM version 2 hello packets. |
| V2 Register | PIM version 2 register packets. |
| V2 Register Stop | PIM version 2 register stop packets. |
| V2 Join Prune | PIM version 2 join and prune packets. |
| V2 Bootstrap | PIM version 2 bootstrap packets. |
| V2 Assert | PIM version 2 assert packets. |
| V2 Graft | PIM version 2 graft packets. |
| V2 Graft Ack | PIM version 2 graft acknowledgment packets. |
| V2 Candidate RP | PIM version 2 candidate RP packets. |
| V2 State Refresh | <p>PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh.</p> <p>State refresh is an extension to PIM-DM. It not supported in Junos OS.</p> |

Table 39: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------------|---|
| V2 DF Election | PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election. |
| V1 Query | PIM version 1 query packets. |
| V1 Register | PIM version 1 register packets. |
| V1 Register Stop | PIM version 1 register stop packets. |
| V1 Join Prune | PIM version 1 join and prune packets. |
| V1 RP Reachability | PIM version 1 RP reachability packets. |
| V1 Assert | PIM version 1 assert packets. |
| V1 Graft | PIM version 1 graft packets. |
| V1 Graft Ack | PIM version 1 graft acknowledgment packets. |
| AutoRP Announce | Auto-RP announce packets. |
| AutoRP Mapping | Auto-RP mapping packets. |
| AutoRP Unknown type | Auto-RP packets with an unknown type. |
| Anycast Register | Auto-RP announce packets. |
| Anycast Register Stop | Auto-RP announce packets. |
| Global Statistics | Summary of PIM statistics for all interfaces. |
| Hello dropped on neighbor policy | Number of hello packets dropped because of a configured neighbor policy. |
| Unknown type | Number of PIM control packets received with an unknown type. |
| V1 Unknown type | Number of PIM version 1 control packets received with an unknown type. |
| Unknown Version | Number of PIM control packets received with an unknown version. The version is not version 1 or version 2. |
| Neighbor unknown | Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet. |
| Bad Length | Number of PIM control packets received for which the packet size does not match the PIM length field in the packet. |

Table 39: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|---------------------------------|---|
| Bad Checksum | Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet. |
| Bad Receive If | Number of PIM control packets received on an interface that does not have PIM configured. |
| Rx Bad Data | Number of PIM control packets received that contain data for TCP Bad register packets. |
| Rx Intf disabled | Number of PIM control packets received on an interface that has PIM disabled. |
| Rx V1 Require V2 | Number of PIM version 1 control packets received on an interface configured for PIM version 2. |
| Rx V2 Require V1 | Number of PIM version 2 control packets received on an interface configured for PIM version 1. |
| Rx Register not RP | Number of PIM register packets received when the routing device is not the RP for the group. |
| Rx Register no route | Number of PIM register packets received when the RP does not have a unicast route back to the source. |
| Rx Register no decap if | Number of PIM register packets received when the RP does not have a de-encapsulation interface. |
| Null Register Timeout | Number of NULL register timeout packets. |
| RP Filtered Source | Number of PIM packets received when the routing device has a source address filter configured for the RP. |
| Rx Unknown Reg Stop | Number of register stop messages received with an unknown type. |
| Rx Join/Prune no state | Number of join and prune messages received for which the routing device has no state. |
| Rx Join/Prune on upstream if | Number of join and prune messages received on the interface used to reach the upstream routing device, toward the RP. |
| Rx Join/Prune for invalid group | Number of join or prune messages received for invalid multicast group addresses. |
| Rx Join/Prune messages dropped | Number of join and prune messages received and dropped. |
| Rx sparse join for dense group | Number of PIM sparse mode join messages received for a group that is configured for dense mode. |

Table 39: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------------|---|
| Rx Graft/Graft Ack no state | Number of graft and graft acknowledgment messages received for which the router or switch has no state. |
| Rx Graft on upstream if | Number of graft messages received on the interface used to reach the upstream routing device, toward the RP. |
| Rx CRP not BSR | Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap. |
| Rx BSR when BSR | Number of BSR messages received in which the PIM message type is Bootstrap. |
| Rx BSR not RPF if | Number of BSR messages received on an interface that is not the RPF interface. |
| Rx unknown hello opt | Number of PIM hello packets received with options that Junos OS does not support. |
| Rx data no state | Number of PIM control packets received for which the routing device has no state for the data type. |
| Rx RP no state | Number of PIM control packets received for which the routing device has no state for the RP. |
| Rx aggregate | Number of PIM aggregate MDT packets received. |
| Rx malformed packet | Number of PIM control packets received with a malformed IP unicast or multicast address family. |
| No RP | Number of PIM control packets received with no RP address. |
| No register encaps if | Number of PIM register packets received when the first-hop routing device does not have an encapsulation interface. |
| No route upstream | Number of PIM control packets received when the routing device does not have a unicast route to the the interface used to reach the upstream routing device, toward the RP. |
| Nexthop Unusable | Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down. |
| RP mismatch | Number of PIM control packets received for which the routing device has an RP mismatch. |
| RP mode mismatch | RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages. |
| RPF neighbor unknown | Number of PIM control packets received for which the routing device has an unknown RPF neighbor for the source. |

Table 39: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|-------------------------------------|--|
| Rx Joins/Prunes filtered | The number of join and prune messages filtered because of configured route filters and source address filters. |
| Tx Joins/Prunes filtered | The number of join and prune messages filtered because of configured route filters and source address filters. |
| Embedded-RP invalid addr | Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains. |
| Embedded-RP limit exceed | Number of times the limit configured with the maximum-rps statement is exceeded. The maximum-rps statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100. |
| Embedded-RP added | <p>Number of packets in which the embedded RP for IPv6 is added.</p> <p>The following receive events trigger extraction of an IPv6 embedded RP address on the routing device:</p> <ul style="list-style-type: none"> • Multicast Listener Discovery (MLD) report for an embedded RP multicast group address • PIM join message with an embedded RP multicast group address • Static embedded RP multicast group address associated with an interface • Packets sent to an embedded RP multicast group address received on the DR <p>An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.</p> |
| Embedded-RP removed | Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature. |
| Rx Register msgs filtering drop | Number of received register messages dropped because of a filter configured for PIM register messages. |
| Tx Register msgs filtering drop | Number of register messages dropped because of a filter configured for PIM register messages. |
| Rx Bidir Join/Prune on non-Bidir if | Error counter for join and prune messages received on non-bidirectional PIM interfaces. |
| Rx Bidir Join/Prune on non-DF if | Error counter for join and prune messages received on non-designated forwarder interfaces. |
| V4 (S,G) Maximum | Maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted. |

Table 39: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------------|---|
| V4 (S,G) Accepted | Number of accepted (S,G) IPv4 multicast routes. |
| V4 (S,G) Threshold | Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device). |
| V4 (S,G) Log Interval | Time (in seconds) between consecutive log messages. |
| V6 (S,G) Maximum | Maximum number of (S,G) IPv6 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted. |
| V6 (S,G) Accepted | Number of accepted (S,G) IPv6 multicast routes. |
| V6 (S,G) Threshold | Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv6 multicast routes accepted by the device). |
| V6 (S,G) Log Interval | Time (in seconds) between consecutive log messages. |
| V4 (grp-prefix, RP) Maximum | Maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted. |
| V4 (grp-prefix, RP) Accepted | Number of accepted group-to-RP IPv4 multicast mappings. |
| V4 (grp-prefix, RP) Threshold | Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device). |
| V4 (grp-prefix, RP) Log Interval | Time (in seconds) between consecutive log messages. |
| V6 (grp-prefix, RP) Maximum | Maximum number of group-to-RP IPv6 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted. |
| V6 (grp-prefix, RP) Accepted | Number of accepted group-to-RP IPv6 multicast mappings. |
| V6 (grp-prefix, RP) Threshold | Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv6 multicast mappings accepted by the device). |
| V6 (grp-prefix, RP) Log Interval | Time (in seconds) between consecutive log messages. |

Table 39: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|---------------------------------------|---|
| V4 Register Maximum | Maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP. |
| V4 Register Accepted | Number of accepted IPv4 PIM registers. |
| V4 Register Threshold | Threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device). |
| V4 Register Log Interval | Time (in seconds) between consecutive log messages. |
| V6 Register Maximum | Maximum number of IPv6 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP. |
| V6 Register Accepted | Number of accepted IPv6 PIM registers. |
| V6 Register Threshold | Threshold at which a warning message is logged (percentage of the maximum number of IPv6 PIM registers accepted by the device). |
| V6 Register Log Interval | Time (in seconds) between consecutive log messages. |
| (*G) Join drop due to SSM range check | PIM join messages that are dropped because the multicast addresses are outside of the SSM address range of 232.0.0.0 through 232.255.255.255. You can extend the accepted SSM address range by configuring the ssm-groups statement. |

Sample Output

show pim statistics

```

user@host> show pim statistics
PIM Message type    Received    Sent    Rx errors
V2 Hello            15          32         0
V2 Register          0          362        0
V2 Register Stop     483          0         0
V2 Join Prune        18          518        0
V2 Bootstrap         0            0         0
V2 Assert            0            0         0
V2 Graft             0            0         0
V2 Graft Ack         0            0         0
V2 Candidate RP      0            0         0
V2 State Refresh     0            0         0
V2 DF Election       0            0         0
V1 Query             0            0         0
V1 Register          0            0         0
V1 Register Stop     0            0         0
V1 Join Prune        0            0         0

```

| | | | |
|-----------------------|---|---|---|
| V1 RP Reachability | 0 | 0 | 0 |
| V1 Assert | 0 | 0 | 0 |
| V1 Graft | 0 | 0 | 0 |
| V1 Graft Ack | 0 | 0 | 0 |
| AutoRP Announce | 0 | 0 | 0 |
| AutoRP Mapping | 0 | 0 | 0 |
| AutoRP Unknown type | 0 | | |
| Anycast Register | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

Global Statistics

| | |
|---|---|
| Hello dropped on neighbor policy | 0 |
| Unknown type | 0 |
| V1 Unknown type | 0 |
| Unknown Version | 0 |
| ipv4 BSR pkt drop due to excessive rate | 0 |
| ipv6 BSR pkt drop due to excessive rate | 0 |
| Neighbor unknown | 0 |
| Bad Length | 0 |
| Bad Checksum | 0 |
| Bad Receive If | 0 |
| Rx Bad Data | 0 |
| Rx Intf disabled | 0 |
| Rx V1 Require V2 | 0 |
| Rx V2 Require V1 | 0 |
| Rx Register not RP | 0 |
| Rx Register no route | 0 |
| Rx Register no decap if | 0 |
| Null Register Timeout | 0 |
| RP Filtered Source | 0 |
| Rx Unknown Reg Stop | 0 |
| Rx Join/Prune no state | 0 |
| Rx Join/Prune on upstream if | 0 |
| Rx Join/Prune for invalid group | 5 |
| Rx Join/Prune messages dropped | 0 |
| Rx sparse join for dense group | 0 |
| Rx Graft/Graft Ack no state | 0 |
| Rx Graft on upstream if | 0 |
| Rx CRP not BSR | 0 |
| Rx BSR when BSR | 0 |
| Rx BSR not RPF if | 0 |
| Rx unknown hello opt | 0 |
| Rx data no state | 0 |
| Rx RP no state | 0 |
| Rx aggregate | 0 |
| Rx malformed packet | 0 |
| Rx illegal TTL | 0 |
| Rx illegal destination address | 0 |
| No RP | 0 |
| No register encap if | 0 |
| No route upstream | 0 |
| Nexthop Unusable | 0 |
| RP mismatch | 0 |
| RP mode mismatch | 0 |
| RPF neighbor unknown | 0 |
| Rx Joins/Prunes filtered | 0 |
| Tx Joins/Prunes filtered | 0 |
| Embedded-RP invalid addr | 0 |
| Embedded-RP limit exceed | 0 |
| Embedded-RP added | 0 |


```

Embedded-RP removed                0
Rx Register msgs filtering drop      0
Tx Register msgs filtering drop      0
Rx Bidir Join/Prune on non-Bidir if  0
Rx Bidir Join/Prune on non-DF if     0
(*,G) Join drop due to SSM range check 0

```

Sample Output

show pim statistics inet interface <interface-name>

```

user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET

```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello | 0 | 4 | 0 |
| V2 Register | 0 | 0 | 0 |
| V2 Register Stop | 0 | 0 | 0 |
| V2 Join Prune | 0 | 0 | 0 |
| V2 Bootstrap | 0 | 0 | 0 |
| V2 Assert | 0 | 0 | 0 |
| V2 Graft | 0 | 0 | 0 |
| V2 Graft Ack | 0 | 0 | 0 |
| V2 Candidate RP | 0 | 0 | 0 |
| V1 Query | 0 | 0 | 0 |
| V1 Register | 0 | 0 | 0 |
| V1 Register Stop | 0 | 0 | 0 |
| V1 Join Prune | 0 | 0 | 0 |
| V1 RP Reachability | 0 | 0 | 0 |
| V1 Assert | 0 | 0 | 0 |
| V1 Graft | 0 | 0 | 0 |
| V1 Graft Ack | 0 | 0 | 0 |
| AutoRP Announce | 0 | 0 | 0 |
| AutoRP Mapping | 0 | 0 | 0 |
| AutoRP Unknown type | 0 | | |
| Anycast Register | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

Sample Output

show pim statistics inet6 interface <interface-name>

```

user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6

```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello | 0 | 4 | 0 |
| V2 Register | 0 | 0 | 0 |
| V2 Register Stop | 0 | 0 | 0 |
| V2 Join Prune | 0 | 0 | 0 |
| V2 Bootstrap | 0 | 0 | 0 |
| V2 Assert | 0 | 0 | 0 |
| V2 Graft | 0 | 0 | 0 |
| V2 Graft Ack | 0 | 0 | 0 |
| V2 Candidate RP | 0 | 0 | 0 |
| Anycast Register | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

show pim statistics instance <instance-name>

```

user@host> show pim statistics instance VPN-A
PIM Message type      Received      Sent  Rx errors
V2 Hello               31           37      0
V2 Register            0            0      0
V2 Register Stop       0            0      0
V2 Join Prune          0           16      0
V2 Bootstrap           0            0      0
V2 Assert              0            0      0
V2 Graft               0            0      0
V2 Graft Ack           0            0      0
V2 Candidate RP        0            0      0
V2 State Refresh       0            0      0
V2 DF Election         0            0      0
V1 Query               0            0      0
V1 Register            0            0      0
V1 Register Stop       0            0      0
V1 Join Prune          0            0      0
V1 RP Reachability     0            0      0
V1 Assert              0            0      0
V1 Graft               0            0      0
V1 Graft Ack           0            0      0
AutoRP Announce        0            0      0
AutoRP Mapping         0            0      0
AutoRP Unknown type    0            0      0
Anycast Register       0            0      0
Anycast Register Stop  0            0      0

```

Global Statistics

```

Hello dropped on neighbor policy      0
Unknown type                          0
V1 Unknown type                       0
Unknown Version                       0
Neighbor unknown                      0
Bad Length                            0
Bad Checksum                          0
Bad Receive If                        0
Rx Bad Data                           0
Rx Intf disabled                      0
Rx V1 Require V2                      0
Rx V2 Require V1                      0
Rx Register not RP                    0
Rx Register no route                  0
Rx Register no decap if               0
Null Register Timeout                 0
RP Filtered Source                    0
Rx Unknown Reg Stop                   0
Rx Join/Prune no state                0
Rx Join/Prune on upstream if          0
Rx Join/Prune for invalid group        0
Rx Join/Prune messages dropped         0
Rx sparse join for dense group         0
Rx Graft/Graft Ack no state           0
Rx Graft on upstream if               0
Rx CRP not BSR                        0
Rx BSR when BSR                       0
Rx BSR not RPF if                     0
Rx unknown hello opt                  0
Rx data no state                      0

```

| | |
|--|-----|
| Rx RP no state | 0 |
| Rx aggregate | 0 |
| Rx malformed packet | 0 |
| Rx illegal TTL | 0 |
| Rx illegal destination address | 0 |
| No RP | 0 |
| No register encap if | 0 |
| No route upstream | 28 |
| Nexthop Unusable | 0 |
| RP mismatch | 0 |
| RP mode mismatch | 0 |
| RPF neighbor unknown | 0 |
| Rx Joins/Prunes filtered | 0 |
| Tx Joins/Prunes filtered | 0 |
| Embedded-RP invalid addr | 0 |
| Embedded-RP limit exceed | 0 |
| Embedded-RP added | 0 |
| Embedded-RP removed | 0 |
| Rx Register msgs filtering drop | 0 |
| Tx Register msgs filtering drop | 0 |
| Rx Bidir Join/Prune on non-Bidir if | 0 |
| Rx Bidir Join/Prune on non-DF if | 0 |
| V4 (S,G) Maximum | 10 |
| V4 (S,G) Accepted | 9 |
| V4 (S,G) Threshold | 80 |
| V4 (S,G) Log Interval | 80 |
| V6 (S,G) Maximum | 8 |
| V6 (S,G) Accepted | 8 |
| V6 (S,G) Threshold | 50 |
| V6 (S,G) Log Interval | 100 |
| V4 (grp-prefix, RP) Maximum | 100 |
| V4 (grp-prefix, RP) Accepted | 5 |
| V4 (grp-prefix, RP) Threshold | 80 |
| V4 (grp-prefix, RP) Log Interval | 10 |
| V6 (grp-prefix, RP) Maximum | 20 |
| V6 (grp-prefix, RP) Accepted | 0 |
| V6 (grp-prefix, RP) Threshold | 90 |
| V6 (grp-prefix, RP) Log Interval | 20 |
| V4 Register Maximum | 100 |
| V4 Register Accepted | 10 |
| V4 Register Threshold | 80 |
| V4 Register Log Interval | 10 |
| V6 Register Maximum | 20 |
| V6 Register Accepted | 0 |
| V6 Register Threshold | 90 |
| V6 Register Log Interval | 20 |
| (*,G) Join drop due to SSM range check | 0 |

Sample Output

show pim statistics interface <interface-name>

```

user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type      Received      Sent  Rx errors
V2 Hello                0             3         0
V2 Register             0             0         0
V2 Register Stop        0             0         0

```

| | | | |
|-----------------------|---|---|---|
| V2 Join Prune | 0 | 0 | 0 |
| V2 Bootstrap | 0 | 0 | 0 |
| V2 Assert | 0 | 0 | 0 |
| V2 Graft | 0 | 0 | 0 |
| V2 Graft Ack | 0 | 0 | 0 |
| V2 Candidate RP | 0 | 0 | 0 |
| V1 Query | 0 | 0 | 0 |
| V1 Register | 0 | 0 | 0 |
| V1 Register Stop | 0 | 0 | 0 |
| V1 Join Prune | 0 | 0 | 0 |
| V1 RP Reachability | 0 | 0 | 0 |
| V1 Assert | 0 | 0 | 0 |
| V1 Graft | 0 | 0 | 0 |
| V1 Graft Ack | 0 | 0 | 0 |
| AutoRP Announce | 0 | 0 | 0 |
| AutoRP Mapping | 0 | 0 | 0 |
| AutoRP Unknown type | 0 | | |
| Anycast Register | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

| PIM Message type | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello | 0 | 3 | 0 |
| V2 Register | 0 | 0 | 0 |
| V2 Register Stop | 0 | 0 | 0 |
| V2 Join Prune | 0 | 0 | 0 |
| V2 Bootstrap | 0 | 0 | 0 |
| V2 Assert | 0 | 0 | 0 |
| V2 Graft | 0 | 0 | 0 |
| V2 Graft Ack | 0 | 0 | 0 |
| V2 Candidate RP | 0 | 0 | 0 |
| Anycast Register | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |