

Multicast Protocols Feature Guide for QFX10000 Switches

Release
15.1x53



Modified: 2016-05-16

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.

Multicast Protocols Feature Guide for QFX10000 Switches

15.1x53

Copyright © 2016, 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	xvii
	Documentation and Release Notes	xvii
	Supported Platforms	xvii
	Using the Examples in This Manual	xvii
	Merging a Full Example	xviii
	Merging a Snippet	xviii
	Documentation Conventions	xix
	Documentation Feedback	xxi
	Requesting Technical Support	xxi
	Self-Help Online Tools and Resources	xxi
	Opening a Case with JTAC	xxii
Chapter 1	Overview	23
	Multicast Overview	23
	Comparing Multicast to Unicast	23
	IP Multicast Uses	25
	IP Multicast Terminology	26
	Reverse-Path Forwarding for Loop Prevention	27
	Shortest-Path Tree for Loop Prevention	27
	Administrative Scoping for Loop Prevention	28
	Multicast Leaf and Branch Terminology	28
	IP Multicast Addressing	28
	Multicast Addresses	29
	Layer 2 Frames and IPv4 Multicast Addresses	29
	Multicast Interface Lists	31
	Multicast Routing Protocols	32
	T Series Router Multicast Performance	35
Part 1	Managing Group Membership	
Chapter 2	Using IGMP	39
	Understanding Group Membership Protocols	39
	Understanding IGMP	40
	Configuring IGMP	42
	Enabling IGMP	43
	Changing the IGMP Version	44
	Modifying the IGMP Host-Query Message Interval	45
	Modifying the IGMP Last-Member Query Interval	46
	Specifying Immediate-Leave Host Removal for IGMP	47
	Filtering Unwanted IGMP Reports at the IGMP Interface Level	48
	Accepting IGMP Messages from Remote Subnetworks	49

	Modifying the IGMP Query Response Interval	50
	Modifying the IGMP Robustness Variable	51
	Limiting the Maximum IGMP Message Rate	52
	Enabling IGMP Static Group Membership	52
	Recording IGMP Join and Leave Events	59
	Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces	60
	Tracing IGMP Protocol Traffic	61
	Disabling IGMP	63
Chapter 3	Using IGMP Snooping	65
	IGMP Snooping Overview	65
	How IGMP Snooping Works	65
	How IGMP Snooping Works with Routed VLAN Interfaces	66
	How Hosts Join and Leave Multicast Groups	66
	IGMP Snooping and Forwarding Interfaces	67
	General Forwarding Rules	67
	Configuring IGMP Snooping	68
	Configuring VLAN-Specific IGMP Snooping Parameters	69
	Example: Configuring IGMP Snooping	70
	Monitoring IGMP Snooping	72
	Verifying the IGMP Snooping Group Timeout Value	73
Chapter 4	Using MLD	75
	Understanding MLD	75
	Examples: Configuring MLD	78
	Understanding MLD	78
	Configuring MLD	81
	Enabling MLD	82
	Modifying the MLD Version	83
	Modifying the MLD Host-Query Message Interval	83
	Modifying the MLD Query Response Interval	84
	Modifying the MLD Last-Member Query Interval	84
	Specifying Immediate-Leave Host Removal for MLD	85
	Filtering Unwanted MLD Reports at the MLD Interface Level	86
	Example: Modifying the MLD Robustness Variable	87
	Limiting the Maximum MLD Message Rate	88
	Enabling MLD Static Group Membership	89
	Create a MLD Static Group Member	89
	Automatically create static groups	90
	Automatically increment group addresses	91
	Specify multicast source address (in SSM mode)	92
	Automatically specify multicast sources	93
	Automatically increment source addresses	94
	Exclude multicast source addresses (in SSM mode)	95
	Example: Recording MLD Join and Leave Events	96
	Configuring the Number of MLD Multicast Group Joins on Logical Interfaces	98
	Disabling MLD	100

Chapter 5	Using MLD Snooping	101
	Understanding MLD Snooping	101
	How MLD Snooping Works	102
	MLD Message Types	103
	How Hosts Join and Leave Multicast Groups	103
	Support for MLDv2 Multicast Sources	104
	MLD Snooping and Forwarding Interfaces	104
	General Forwarding Rules	105
	Examples of MLD Snooping Multicast Forwarding	105
	Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts	105
	Scenario 2: Switch Forwarding Multicast Traffic to Another Switch	106
	Scenario 3: Switch Connected to Hosts Only (No MLD Querier)	107
	Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs	108
	Configuring MLD Snooping on a VLAN (CLI Procedure)	109
	Enabling or Disabling MLD Snooping on VLANs	110
	Configuring the MLD Version	111
	Enabling Immediate Leave	111
	Configuring an Interface as a Multicast-Router Interface	112
	Configuring Static Group Membership on an Interface	113
	Changing the Timer and Counter Values	114
	Example: Configuring MLD Snooping	115
	Verifying MLD Snooping	118
	Verifying MLD Snooping Memberships	119
	Verifying MLD Snooping Interfaces	119
	Viewing MLD Snooping Statistics	120
	Viewing MLD Snooping Routing Information	121
Part 2	Configuring PIM	
Chapter 6	Using PIM Basic Features	125
	PIM Overview	125
	Basic PIM Network Components	127
	PIM on Aggregated Interfaces	128
	Changing the PIM Version	128
	Modifying the PIM Hello Interval	128
	Preserving Multicast Performance by Disabling Response to the ping Utility	129
	Configuring PIM Trace Options	130
	Configuring Interface Priority for PIM Designated Router Selection	132
	Configuring PIM Designated Router Election on Point-to-Point Links	133
	Configuring BFD for PIM	134
	Configuring BFD Authentication for PIM	135
	Configuring BFD Authentication Parameters	136
	Viewing Authentication Information for BFD Sessions	137
	Disabling PIM	139
	Disabling the PIM Protocol	139
	Disabling PIM on an Interface	140
	Disabling PIM for a Family	140

	Disabling PIM for a Rendezvous Point	141
Chapter 7	Using PIM Sparse Mode	143
	Understanding PIM Sparse Mode	143
	Rendezvous Point	145
	RP Mapping Options	145
	Designated Router	146
	Enabling PIM Sparse Mode	146
	Configuring PIM Join Load Balancing	147
	Modifying the Join State Timeout	151
	Example: Enabling Join Suppression	151
Chapter 8	Using PIM Dense Mode and PIM Sparse-Dense Mode	157
	Understanding PIM Dense Mode	157
	Understanding PIM Sparse-Dense Mode	159
	Mixing PIM Sparse and Dense Modes	159
	Configuring PIM Dense Mode Properties	160
	Configuring PIM Sparse-Dense Mode Properties	161
Chapter 9	Using Source-Specific Multicast	163
	Source-Specific Multicast Groups Overview	163
	Understanding PIM Source-Specific Mode	164
	PIM SSM	165
	Example: Configuring PIM SSM on a Network	167
	Example: Configuring an SSM-Only Domain	169
	Example: Configuring SSM Mapping	169
	Example: Configuring Source-Specific Multicast Groups with Any-Source Override	172
	Example: Configuring SSM Maps for Different Groups to Different Sources	175
	Multiple SSM Maps and Groups for Interfaces	175
	Example: Configuring Multiple SSM Maps Per Interface	175
Chapter 10	Using Static RP	181
	Understanding Static RP	181
	Configuring Local PIM RPs	181
	Configuring the Static PIM RP Address on the Non-RP Routing Device	183
Chapter 11	Using Anycast RP	185
	Understanding RP Mapping with Anycast RP	185
	Example: Configuring PIM Anycast With or Without MSDP	186
	Configuring a PIM Anycast RP Router with MSDP	189
	Configuring a PIM Anycast RP Router Using Only PIM	190
	Configuring All PIM Anycast Non-RP Routers	191
	Example: Configuring Multiple RPs in a Domain with Anycast RP	192
Chapter 12	Using Auto-RP	195
	Understanding PIM Auto-RP	195
	Configuring PIM Auto-RP	195

Chapter 13	Using PIM Bootstrap Router	201
	Understanding the PIM Bootstrap Router	201
	Configuring PIM Bootstrap Properties for IPv4 or IPv6	201
	Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain	203
	Example: Configuring PIM BSR Filters	203
Chapter 14	Using PIM Filtering	205
	Understanding Multicast Message Filters	205
	Filtering MAC Addresses	206
	Filtering RP and DR Register Messages	206
	Configuring Interface-Level PIM Neighbor Policies	207
	Filtering Outgoing PIM Join Messages	208
	Filtering Incoming PIM Join Messages	209
	Configuring Register Message Filters on a PIM RP and DR	211
Chapter 15	Using PIM RPT and SPT Cutover	213
	Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees	213
	Building an RPT Between the RP and Receivers	214
	PIM Sparse Mode Source Registration	215
	Multicast Shortest-Path Tree	218
	SPT Cutover	219
	SPT Cutover Control	222
	Example: Configuring the PIM Assert Timeout	222
	Example: Configuring the PIM SPT Threshold Policy	224
Part 3	Configuring MSDP	
Chapter 16	Using MSDP	231
	Understanding MSDP	231
	Configuring MSDP	232
	Filtering MSDP SA Messages	234
	Tracing MSDP Protocol Traffic	234
	Configuring the Interface to Accept Traffic from a Remote Source	236
	Example: Configuring MSDP	237
	Example: Configuring MSDP with Active Source Limits and Mesh Groups	238
	Example: Configuring PIM Anycast With or Without MSDP	244
	Configuring a PIM Anycast RP Router with MSDP	247
Part 4	Configuration Statements and Operational Commands	
Chapter 17	Configuration Statements (IGMP)	251
	accounting (Protocols IGMP)	252
	accounting (Protocols IGMP Interface)	252
	asm-override-ssm	253
	disable (Protocols IGMP)	253
	exclude (Protocols IGMP)	254
	group (Protocols IGMP)	255
	group-count	256

	group-increment (Protocols IGMP)	256
	group-limit (Protocols IGMP)	257
	group-policy (Protocols IGMP)	257
	igmp	258
	immediate-leave (Protocols IGMP)	260
	interface (Protocols IGMP)	261
	maximum-transmit-rate (Protocols IGMP)	262
	oif-map (IGMP Interface)	262
	passive (IGMP)	263
	promiscuous-mode (Protocols IGMP)	264
	query-interval (Protocols IGMP)	265
	query-last-member-interval (Protocols IGMP)	266
	query-response-interval (Protocols IGMP)	267
	robust-count (Protocols IGMP)	268
	source (Protocols IGMP)	269
	source-count (Protocols IGMP)	270
	source-increment (Protocols IGMP)	271
	static (Protocols IGMP)	272
	traceoptions (Protocols IGMP)	273
	version (Protocols IGMP)	275
Chapter 18	Configuration Statements (IGMP Snooping)	277
	all	278
	data-forwarding	278
	disable (IGMP Snooping)	279
	group (IGMP Snooping)	279
	group-limit (IGMP and MLD Snooping)	280
	host-only-interface	281
	igmp-querier	281
	igmp-snooping	282
	immediate-leave (Bridge Domains)	283
	interface (Bridge Domains)	284
	interface (IGMP Snooping)	285
	l2-querier	285
	multicast-router-interface (IGMP Snooping)	286
	query-interval (Bridge Domains)	287
	query-last-member-interval (Bridge Domains)	288
	query-response-interval (Bridge Domains)	289
	receiver	290
	robust-count (IGMP Snooping)	290
	source-address	291
	src-address (IGMP Querier)	292
	source-vlans	292
	static (IGMP Snooping)	293
	traceoptions (IGMP Snooping)	294
	version (IGMP Snooping)	296
	vlan (IGMP Snooping)	297

Chapter 19	Configuration Statements (MLD Snooping)	299
	all	300
	data-forwarding	300
	disable (IGMP Snooping)	301
	group (IGMP Snooping)	301
	group-limit (IGMP and MLD Snooping)	302
	host-only-interface	303
	igmp-querier	303
	igmp-snooping	304
	immediate-leave (Bridge Domains)	305
	interface (Bridge Domains)	306
	interface (IGMP Snooping)	307
	l2-querier	307
	mld-snooping	308
	multicast-router-interface (IGMP Snooping)	309
	query-interval (Bridge Domains)	310
	query-last-member-interval (Bridge Domains)	311
	query-response-interval (Bridge Domains)	312
	receiver	313
	robust-count (IGMP Snooping)	313
	source-address	314
	src-address (IGMP Querier)	315
	source-vlans	315
	static (IGMP Snooping)	316
	traceoptions (IGMP Snooping)	317
	version (IGMP Snooping)	319
	vlan (IGMP Snooping)	320
Chapter 20	Configuration Statements (PIM)	321
	address (Anycast RPs)	323
	address (Local RPs)	324
	address (Static RPs)	325
	algorithm	326
	anycast-pim	327
	assert-timeout	328
	authentication	329
	auto-rp	330
	bfd-liveness-detection	331
	bootstrap	332
	bootstrap-export	333
	bootstrap-import	334
	bootstrap-priority	335
	dense-groups	336
	detection-time (BFD for PIM)	337
	disable (PIM)	338
	dr-election-on-p2p	339
	dr-register-policy	339
	embedded-rp	340
	export (Bootstrap)	341

export (Protocols PIM)	342
family (Bootstrap)	343
family (Protocols PIM)	344
family (Local RP)	345
group (RPF Selection)	346
group-ranges	347
hello-interval	348
hold-time (Protocols PIM)	349
import (Protocols PIM Bootstrap)	350
import (Protocols PIM)	351
infinity	352
interface	353
join-load-balance	354
join-prune-timeout	355
key-chain	355
local	356
local-address (Protocols PIM)	357
loose-check	358
mapping-agent-election	359
maximum-rps	360
minimum-interval (PIM BFD Liveness Detection)	361
minimum-interval (PIM BFD Transmit Interval)	362
minimum-receive-interval	363
mode (Protocols PIM)	363
multiplier	364
neighbor-policy	364
next-hop (PIM RPF Selection)	365
no-adaptation (PIM BFD Liveness Detection)	365
override-interval	366
pim	367
prefix-list (PIM RPF Selection)	370
priority (Bootstrap)	371
priority (PIM Interfaces)	372
priority (PIM RPs)	373
propagation-delay	374
register-probe-time	375
reset-tracking-bit	376
rib-group (Protocols PIM)	377
rp	378
rp-register-policy	380
rp-set	381
rpf-selection	382
source (PIM RPF Selection)	383
spt-threshold	384
static (Protocols PIM)	385
threshold (PIM BFD Detection Time)	386
threshold (PIM BFD Transmit Interval)	387
transmit-interval (PIM BFD Liveness Detection)	388
traceoptions (Protocols PIM)	389

	version (BFD)	392
	version (PIM)	393
	wildcard-source (PIM RPF Selection)	394
Chapter 21	Configuration Statements (Source-Specific Multicast)	395
	asm-override-ssm	395
	policy (SSM Maps)	396
	ssm-groups	397
	ssm-map (Protocols IGMP)	398
	ssm-map (Routing Options Multicast)	399
	ssm-map-policy (IGMP)	400
Chapter 22	Configuration Statements (MSDP)	401
	active-source-limit	402
	authentication-key	403
	data-encapsulation	404
	default-peer	405
	disable (Protocols MSDP)	406
	export (Protocols MSDP)	407
	group (Protocols MSDP)	408
	import (Protocols MSDP)	409
	local-address (Protocols MSDP)	410
	maximum (MSDP Active Source Messages)	411
	mode (Protocols MSDP)	412
	msdp	413
	peer (Protocols MSDP)	415
	rib-group (Protocols MSDP)	416
	source (Protocols MSDP)	417
	threshold (MSDP Active Source Messages)	418
	traceoptions (Protocols MSDP)	419
Chapter 23	Operational Commands (IGMP)	423
	clear igmp membership	424
	clear igmp statistics	427
	show igmp group	429
	show configuration protocols igmp	433
	show igmp interface	435
	show igmp statistics	439
	show system statistics igmp	442
Chapter 24	Operational Commands (IGMP Snooping)	447
	clear igmp-snooping membership	448
	clear igmp-snooping statistics	449
	show igmp-snooping membership	450
	show igmp-snooping route	453
	show igmp-snooping statistics	455
	show igmp-snooping vlans	457
Chapter 25	Operational Commands (PIM)	459
	clear multicast bandwidth-admission	461
	clear multicast scope	463

	clear multicast sessions	464
	clear multicast statistics	465
	clear pim join	466
	clear pim register	468
	clear pim statistics	470
	mtrace	473
	mtrace from-source	476
	mtrace monitor	479
	mtrace to-gateway	481
	show multicast flow-map	484
	show multicast interface	486
	show multicast mrinfo	488
	show multicast next-hops	490
	show multicast pim-to-igmp-proxy	493
	show multicast pim-to-mld-proxy	495
	show multicast route	497
	show multicast rpf	503
	show multicast scope	507
	show multicast sessions	509
	show multicast usage	512
	show pim bootstrap	515
	show pim interfaces	517
	show pim join	520
	show pim neighbors	541
	show pim rps	545
	show pim source	552
	show pim statistics	555
Chapter 26	Operational Commands (MSDP)	565
	clear msdp cache	566
	clear msdp statistics	567
	show msdp	568
	show msdp source	570
	show msdp source-active	572
	show msdp statistics	575
	test msdp	579
Part 2	Index	
	Index	583

List of Figures

Chapter 1	Overview	23
	Figure 1: Multicast Terminology in an IP Network	27
	Figure 2: Converting MAC Addresses to Multicast Addresses	31
Part 1	Managing Group Membership	
Chapter 4	Using MLD	75
	Figure 3: Routing Devices Start Up on a Subnet	76
	Figure 4: Querier Routing Device Is Determined	76
	Figure 5: General Query Message Is Issued	77
	Figure 6: Reports Are Received by the Querier Routing Device	77
	Figure 7: Host Has No Interested Receivers and Sends a Done Message to Routing Device	77
	Figure 8: Host Address Timer Expires and Address Is Removed from Multicast Address List	78
	Figure 9: Routing Devices Start Up on a Subnet	79
	Figure 10: Querier Routing Device Is Determined	80
	Figure 11: General Query Message Is Issued	80
	Figure 12: Reports Are Received by the Querier Routing Device	80
	Figure 13: Host Has No Interested Receivers and Sends a Done Message to Routing Device	81
	Figure 14: Host Address Timer Expires and Address Is Removed from Multicast Address List	81
Chapter 5	Using MLD Snooping	101
	Figure 15: Multicast Traffic Flow with MLD Snooping Enabled	102
	Figure 16: Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts	106
	Figure 17: Scenario 2: Switch Forwarding Multicast Traffic to Another Switch . . .	107
	Figure 18: Scenario 3: Switch Connected to Hosts Only (No MLD Querier)	108
	Figure 19: Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs	109
	Figure 20: MLD Snooping Topology Example	116
Part 2	Configuring PIM	
Chapter 7	Using PIM Sparse Mode	143
	Figure 21: Rendezvous Point As Part of the RPT and SPT	145
	Figure 22: Join Suppression	153
Chapter 8	Using PIM Dense Mode and PIM Sparse-Dense Mode	157
	Figure 23: Multicast Traffic Flooded from the Source Using PIM Dense Mode . . .	158

	Figure 24: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic	159
Chapter 9	Using Source-Specific Multicast	163
	Figure 25: Receiver Announces Desire to Join Group G and Source S	166
	Figure 26: Router 3 (Last-Hop Router) Joins the Source Tree	166
	Figure 27: (S,G) State Is Built Between the Source and the Receiver	166
	Figure 28: Network on Which to Configure PIM SSM	167
	Figure 29: Receiver Sends Messages to Join Group G and Source S	172
	Figure 30: Router 3 (Last-Hop Router) Joins the Source Tree	173
	Figure 31: (S,G) State Is Built Between the Source and the Receiver	173
	Figure 32: Simple RPF Topology	173
Chapter 15	Using PIM RPT and SPT Cutover	213
	Figure 33: Building an RPT Between the RP and the Receiver	215
	Figure 34: PIM Register Message and PIM Join Message Exchanged	216
	Figure 35: Traffic Sent from the Source to the RP Router	217
	Figure 36: Traffic Sent from the RP Router Toward the Receiver	217
	Figure 37: Receiver DR Sends a PIM Join Message to the Source	219
	Figure 38: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router	220
	Figure 39: RP Router Receives PIM Prune Message	220
	Figure 40: RP Router Sends a PIM Prune Message to the Source DR	221
	Figure 41: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router	221
	Figure 42: PIM Assert Topology	223
Part 3	Configuring MSDP	
Chapter 16	Using MSDP	231
	Figure 43: Accepting Multicast Traffic from a Remote Source	236
	Figure 44: Source-Active Message Flooding	241

List of Tables

	About the Documentation	xvii
	Table 1: Notice Icons	xix
	Table 2: Text and Syntax Conventions	xix
Chapter 1	Overview	23
	Table 3: Multicast Routing Protocols Compared	34
Part 1	Managing Group Membership	
Chapter 2	Using IGMP	39
	Table 4: IGMP Event Messages	59
Chapter 3	Using IGMP Snooping	65
	Table 5: Components of the IGMP Snooping Topology	70
	Table 6: Summary of IGMP Snooping Output Fields	72
Chapter 4	Using MLD	75
	Table 7: MLD Event Messages	96
Part 2	Configuring PIM	
Chapter 9	Using Source-Specific Multicast	163
	Table 8: ASM and SSM Terminology	165
Chapter 12	Using Auto-RP	195
	Table 9: Local RP and Auto-RP Message Types	196
Chapter 14	Using PIM Filtering	205
	Table 10: PIM Join Filter Match Conditions	210
Part 3	Configuring MSDP	
Chapter 16	Using MSDP	231
	Table 11: Source-Active Message Flooding Explanation	240
Part 4	Configuration Statements and Operational Commands	
Chapter 23	Operational Commands (IGMP)	423
	Table 12: show igmp group Output Fields	429
	Table 13: show igmp group Output Fields	433
	Table 14: show igmp interface Output Fields	435
	Table 15: show igmp statistics Output Fields	439
Chapter 24	Operational Commands (IGMP Snooping)	447

	Table 16: show igmp-snooping membership Output Fields	450
	Table 17: show igmp-snooping route Output Fields	453
	Table 18: show igmp-snooping statistics Output Fields	455
	Table 19: show igmp-snooping vlans Output Fields	457
Chapter 25	Operational Commands (PIM)	459
	Table 20: mtrace Output Fields	473
	Table 21: mtrace from-source Output Fields	477
	Table 22: mtrace monitor Output Fields	479
	Table 23: mtrace to-gateway Output Fields	482
	Table 24: show multicast flow-map Output Fields	484
	Table 25: show multicast interface Output Fields	486
	Table 26: show multicast minfo Output Fields	488
	Table 27: show multicast next-hops Output Fields	491
	Table 28: show multicast pim-to-igmp-proxy Output Fields	493
	Table 29: show multicast pim-to-mld-proxy Output Fields	495
	Table 30: show multicast route Output Fields	498
	Table 31: show multicast rpf Output Fields	504
	Table 32: show multicast scope Output Fields	507
	Table 33: show multicast sessions Output Fields	509
	Table 34: show multicast usage Output Fields	513
	Table 35: show pim bootstrap Output Fields	515
	Table 36: show pim interfaces Output Fields	517
	Table 37: show pim join Output Fields	522
	Table 38: show pim neighbors Output Fields	542
	Table 39: show pim rps Output Fields	546
	Table 40: show pim source Output Fields	553
	Table 41: show pim statistics Output Fields	556
Chapter 26	Operational Commands (MSDP)	565
	Table 42: show msdp Output Fields	568
	Table 43: show msdp source Output Fields	571
	Table 44: show msdp source-active Output Fields	573
	Table 45: show msdp statistics Output Fields	575

About the Documentation

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

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:

- [QFX 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 xix 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 xix defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

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

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

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<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 domain-name
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 metric>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (string1 string2 string3)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop address; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<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 of the Juniper Networks TechLibrary 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 <http://www.juniper.net/techpubs/feedback/>.
- 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>.

CHAPTER 1

Overview

- [Multicast Overview on page 23](#)

Multicast Overview

IP has three fundamental types of addresses: unicast, broadcast, and multicast. A *unicast address* is used to send a packet to a single destination. A *broadcast address* is used to send a datagram to an entire subnetwork. A *multicast address* is used to send a datagram to a set of hosts that can be on different subnetworks and that are configured as members of a multicast group.

A multicast datagram is delivered to destination group members with the same best-effort reliability as a standard unicast IP datagram. This means that multicast datagrams are not guaranteed to reach all members of a group or to arrive in the same order in which they were transmitted. The only difference between a multicast IP packet and a unicast IP packet is the presence of a group address in the IP header destination address field. Multicast addresses use the Class D address format.



NOTE: On all SRX Series devices, reordering is not supported for multicast fragments. Reordering of unicast fragments is supported.

Individual hosts can join or leave a multicast group at any time. There are no restrictions on the physical location or the number of members in a multicast group. A host can be a member of more than one multicast group at any time. A host does not have to belong to a group to send packets to members of a group.

Routers use a group membership protocol to learn about the presence of group members on directly attached subnetworks. When a host joins a multicast group, it transmits a group membership protocol message for the group or groups that it wants to receive and sets its IP process and network interface card to receive frames addressed to the multicast group.

Comparing Multicast to Unicast

The Junos[®] operating system (Junos OS) routing protocol process supports a wide variety of routing protocols. These routing protocols carry network information among routing devices not only for *unicast* traffic streams sent between one pair of clients and servers,

but also for *multicast* traffic streams containing video, audio, or both, between a single server source and many client receivers. The routing protocols used for multicast differ in many key ways from unicast routing protocols.

Information is delivered over a network by three basic methods: unicast, broadcast, and multicast.

The differences among unicast, broadcast, and multicast can be summarized as follows:

- Unicast: One-to-one, from one source to one destination.
- Broadcast: One-to-all, from one source to all possible destinations.
- Multicast: One-to-many, from one source to multiple destinations expressing an interest in receiving the traffic.



NOTE: This list does not include a special category for many-to-many applications, such as online gaming or videoconferencing, where there are many sources for the same receiver and where receivers often double as sources. Many-to-many is a service model that repeatedly employs one-to-many multicast and therefore requires no unique protocol. The original multicast specification, RFC 1112, supports both the any-source multicast (ASM) many-to-many model and the source-specific multicast (SSM) one-to-many model.

With unicast traffic, many streams of IP packets that travel across networks flow from a single source, such as a website server, to a single destination such as a client PC. Unicast traffic is still the most common form of information transfer on networks.

Broadcast traffic flows from a single source to all possible destinations reachable on the network, which is usually a LAN. Broadcasting is the easiest way to make sure traffic reaches its destinations.

Television networks use broadcasting to distribute video and audio. Even if the television network is a cable television (CATV) system, the source signal reaches all possible destinations, which is the main reason that some channels' content is scrambled. Broadcasting is not feasible on the Internet because of the enormous amount of unnecessary information that would constantly arrive at each end user's device, the complexities and impact of scrambling, and related privacy issues.

Multicast traffic lies between the extremes of unicast (one source, one destination) and broadcast (one source, all destinations). Multicast is a "one source, many destinations" method of traffic distribution, meaning only the destinations that explicitly indicate their need to receive the information from a particular source receive the traffic stream.

On an IP network, because destinations (clients) do not often communicate directly with sources (servers), the routing devices between source and destination must be able to determine the topology of the network from the unicast or multicast perspective to avoid routing traffic haphazardly. Multicast routing devices replicate packets received on one input interface and send the copies out on multiple output interfaces.

In IP multicast, the source and destination are almost always hosts and not routing devices. Multicast routing devices distribute the multicast traffic across the network from source to destinations. The multicast routing device must find multicast sources on the network, send out copies of packets on several interfaces, prevent routing loops, connect interested destinations with the proper source, and keep the flow of unwanted packets to a minimum. Standard multicast routing protocols provide most of these capabilities, but some router architectures cannot send multiple copies of packets and so do not support multicasting directly.

IP Multicast Uses

Multicast allows an IP network to support more than just the unicast model of data delivery that prevailed in the early stages of the Internet. Multicast, originally defined as a host extension in RFC 1112 in 1989, provides an efficient method for delivering traffic flows that can be characterized as one-to-many or many-to-many.

Unicast traffic is not strictly limited to data applications. Telephone conversations, wireless or not, contain digital audio samples and might contain digital photographs or even video and still flow from a single source to a single destination. In the same way, multicast traffic is not strictly limited to multimedia applications. In some data applications, the flow of traffic is from a single source to many destinations that require the packets, as in a news or stock ticker service delivered to many PCs. For this reason, the term *receiver* is preferred to *listener* for multicast destinations, although both terms are common.

Network applications that can function with unicast but are better suited for multicast include collaborative groupware, teleconferencing, periodic or “push” data delivery (stock quotes, sports scores, magazines, newspapers, and advertisements), server or website replication, and distributed interactive simulation (DIS) such as war simulations or virtual reality. Any IP network concerned with reducing network resource overhead for one-to-many or many-to-many data or multimedia applications with multiple receivers benefits from multicast.

If unicast were employed by radio or news ticker services, each radio or PC would have to have a separate traffic session for each listener or viewer at a PC (this is actually the method for some Web-based services). The processing load and bandwidth consumed by the server would increase linearly as more people “tune in” to the server. This is extremely inefficient when dealing with the global scale of the Internet. Unicast places the burden of packet duplication on the server and consumes more and more backbone bandwidth as the number of users grows.

If broadcast were employed instead, the source could generate a single IP packet stream using a broadcast destination address. Although broadcast eliminates the server packet duplication issue, this is not a good solution for IP because IP broadcasts can be sent only to a single subnetwork, and IP routing devices normally isolate IP subnetworks on separate interfaces. Even if an IP packet stream could be addressed to literally go everywhere, and there were no need to “tune” to any source at all, broadcast would be extremely inefficient because of the bandwidth strain and need for uninterested hosts to discard large numbers of packets. Broadcast places the burden of packet rejection on each host and consumes the maximum amount of backbone bandwidth.


For radio station or news ticker traffic, multicast provides the most efficient and effective outcome, with none of the drawbacks and all of the advantages of the other methods. A single source of multicast packets finds its way to every *interested* receiver. As with broadcast, the transmitting host generates only a single stream of IP packets, so the load remains constant whether there is one receiver or one million. The network routing devices replicate the packets and deliver the packets to the proper receivers, but only the replication role is a new one for routing devices. The links leading to subnets consisting of entirely uninterested receivers carry no multicast traffic. Multicast minimizes the burden placed on sender, network, and receiver.

IP Multicast Terminology

Multicast has its own particular set of terms and acronyms that apply to IP multicast routing devices and networks. [Figure 1 on page 27](#) depicts some of the terms commonly used in an IP multicast network.

In a multicast network, the key component is the *routing device*, which is able to replicate packets and is therefore multicast-capable. The routing devices in the IP multicast network, which has exactly the same topology as the unicast network it is based on, use a *multicast routing protocol* to build a *distribution tree* that connects receivers (preferred to the multimedia implications of listeners, but listeners is also used) to *sources*. In multicast terminology, the distribution tree is *rooted at the source* (the root of the distribution tree is the source). The interface on the routing device leading toward the source is the *upstream* interface, although the less precise terms *incoming* or *inbound* interface are used as well. To keep bandwidth use to a minimum, it is best for only one upstream interface on the routing device to receive multicast packets. The interface on the routing device leading toward the receivers is the *downstream* interface, although the less precise terms *outgoing* or *outbound* interface are used as well. There can be 0 to $N-1$ downstream interfaces on a routing device, where N is the number of logical interfaces on the routing device. To prevent looping, the upstream interface must never receive copies of downstream multicast packets.

Figure 1: Multicast Terminology in an IP Network



Routing loops are disastrous in multicast networks because of the risk of repeatedly replicated packets. One of the complexities of modern multicast routing protocols is the need to avoid routing loops, packet by packet, much more rigorously than in unicast routing protocols.

Reverse-Path Forwarding for Loop Prevention

The routing device's multicast forwarding state runs more logically based on the reverse path, from the receiver back to the root of the distribution tree. In RPF, every multicast packet received must pass an RPF check before it can be replicated or forwarded on any interface. When it receives a multicast packet on an interface, the routing device verifies that the *source* address in the multicast IP packet is the *destination* address for a unicast IP packet back to the source.

If the outgoing interface found in the unicast routing table is the same interface that the multicast packet was received on, the packet passes the RPF check. Multicast packets that fail the RPF check are dropped, because the incoming interface is not on the shortest path back to the source. routing devices can build and maintain separate tables for RPF purposes.

Shortest-Path Tree for Loop Prevention

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT), but this path can be long if the source is at the periphery of the network. Providing a *shared tree* on the backbone as the distribution tree locates the multicast source more centrally in the network. Shared distribution trees with roots in the core network are created and maintained by a multicast routing device operating as a rendezvous point (RP), a feature of sparse mode multicast protocols.

Administrative Scoping for Loop Prevention

Scoping limits the routing devices and interfaces that can forward a multicast packet. Multicast scoping is *administrative* in the sense that a range of multicast addresses is reserved for scoping purposes, as described in RFC 2365, *Administratively Scoped IP Multicast*. routing devices at the boundary must filter multicast packets and ensure that packets do not stray beyond the established limit.

Multicast Leaf and Branch Terminology

Each subnetwork with hosts on the routing device that has at least one interested receiver is a *leaf* on the distribution tree. routing devices can have multiple leaves on different interfaces and must send a copy of the IP multicast packet out on each interface with a leaf. When a new leaf subnetwork is added to the tree (that is, the interface to the host subnetwork previously received no copies of the multicast packets), a new *branch* is built, the leaf is joined to the tree, and replicated packets are sent out on the interface. The number of leaves on a particular interface does not affect the routing device. The action is the same for one leaf or a hundred.



NOTE: On Juniper Networks security devices, if the maximum number of leaves on a multicast distribution tree is exceeded, multicast sessions are created up to the maximum number of leaves, and any multicast sessions that exceed the maximum number of leaves are ignored. The maximum number of leaves on a multicast distribution tree is device specific.

When a branch contains no leaves because there are no interested hosts on the routing device interface leading to that IP subnetwork, the branch is *pruned* from the distribution tree, and no multicast packets are sent out that interface. Packets are replicated and sent out multiple interfaces only where the distribution tree branches at a routing device, and no link ever carries a duplicate flow of packets.

Collections of hosts all receiving the same stream of IP packets, usually from the same multicast source, are called *groups*. In IP multicast networks, traffic is delivered to multicast groups based on an IP multicast address, or *group address*. The groups determine the location of the leaves, and the leaves determine the branches on the multicast network.

IP Multicast Addressing

Multicast uses the Class D IP address range (224.0.0.0 through 239.255.255.255). Class D addresses are commonly referred to as *multicast addresses* because the entire classful address concept is obsolete. Multicast addresses can never appear as the source address in an IP packet and can only be the destination of a packet.

Multicast addresses usually have a prefix length of /32, although other prefix lengths are allowed. Multicast addresses represent logical groupings of receivers and not physical collections of devices. Blocks of multicast addresses can still be described in terms of prefix length in traditional notation, but only for convenience. For example, the multicast

address range from 232.0.0.0 through 232.255.255.255 can be written as 232.0.0.0/8 or 232/8.

Internet service providers (ISPs) do not typically allocate multicast addresses to their customers because multicast addresses relate to content, not to physical devices. Receivers are not assigned their own multicast addresses, but need to know the multicast address of the content. Sources need to be assigned multicast addresses only to produce the content, not to identify their place in the network. Every source and receiver still needs an ordinary, unicast IP address.

Multicast addressing most often references the receivers, and the source of multicast content is usually not even a member of the multicast group for which it produces content. If the source needs to monitor the packets it produces, monitoring can be done locally, and there is no need to make the packets traverse the network.

Many applications have been assigned a range of multicast addresses for their own use. These applications assign multicast addresses to sessions created by that application. You do not usually need to statically assign a multicast address, but you can do so.

Multicast Addresses

Multicast host group addresses are defined to be the IP addresses whose high-order four bits are 1110, giving an address range from 224.0.0.0 through 239.255.255.255, or simply 224.0.0.0/4. (These addresses also are referred to as Class D addresses.)

The Internet Assigned Numbers Authority (IANA) maintains a list of registered IP multicast groups. The base address 224.0.0.0 is reserved and cannot be assigned to any group. The block of multicast addresses from 224.0.0.1 through 224.0.0.255 is reserved for local wire use. Groups in this range are assigned for various uses, including routing protocols and local discovery mechanisms.

The range from 239.0.0.0 through 239.255.255.255 is reserved for administratively scoped addresses. Because packets addressed to administratively scoped multicast addresses do not cross configured administrative boundaries, and because administratively scoped multicast addresses are locally assigned, these addresses do not need to be unique across administrative boundaries.

Layer 2 Frames and IPv4 Multicast Addresses

Multicasting on a LAN is a good place to start an investigation of multicasting at Layer 2. At Layer 2, multicast deals with media access control (MAC) frames and addresses instead of IPv4 or IPv6 packets and addresses. Consider a single LAN, without routing devices, with a multicast source sending to a certain group. The rest of the hosts are receivers interested in the multicast group's content. So the multicast source host generates packets with its unicast IP address as the source, and the multicast group address as the destination.

Which MAC addresses are used on the frame containing this packet? The packet source address—the unicast IP address of the host originating the multicast content—translates easily and directly to the MAC address of the source. But what about the packet's destination address? This is the IP multicast group address. Which destination MAC address for the frame corresponds to the packet's multicast group address?

One option is for LANs simply to use the LAN broadcast MAC address, which guarantees that the frame is processed by every station on the LAN. However, this procedure defeats the whole purpose of multicast, which is to limit the circulation of packets and frames to interested hosts. Also, hosts might have access to many multicast groups, which multiplies the amount of traffic to noninterested destinations. Broadcasting frames at the LAN level to support multicast groups makes no sense.

However, there is an easy way to effectively use Layer 2 frames for multicast purposes. The MAC address has a bit that is set to 0 for unicast (the LAN term is *individual address*) and set to 1 to indicate that this is a multicast address. Some of these addresses are reserved for multicast groups of specific vendors or MAC-level protocols. Internet multicast applications use the range 0x01-00-5E-00-00-00 to 0x01-00-5E-FF-FF-FF. Multicast receivers (hosts running TCP/IP) listen for frames with one of these addresses when the application joins a multicast group. The host stops listening when the application terminates or the host leaves the group at the packet layer (Layer 3).

This means that 3 bytes, or 24 bits, are available to map IPv4 multicast addresses at Layer 3 to MAC multicast addresses at Layer 2. However, all IPv4 addresses, including multicast addresses, are 32 bits long, leaving 8 IP address bits left over. Which method of mapping IPv4 multicast addresses to MAC multicast addresses minimizes the chance of “collisions” (that is, two different IP multicast groups at the packet layer mapping to the same MAC multicast address at the frame layer)?

First, it is important to realize that all IPv4 multicast addresses begin with the same 4 bits (**1110**), so there are really only 4 bits of concern, not 8. A LAN must not drop the last bits of the IPv4 address because these are almost guaranteed to be host bits, depending on the subnet mask. But the high-order bits, the leftmost address bits, are almost always network bits, and there is only one LAN (for now).

One other bit of the remaining 24 MAC address bits is reserved (an initial **0** indicates an Internet multicast address), so the 5 bits following the initial **1110** in the IPv4 address are dropped. The 23 remaining bits are mapped, one for one, into the last 23 bits of the MAC address. An example of this process is shown in [Figure 2 on page 31](#).

Figure 2: Converting MAC Addresses to Multicast Addresses

1	IPv4 header multicast destination address	232.	224.	202.	181
	Written in hexadecimal	E8	E0	CA	B5
	Written in binary	1110 1000 1	110 0000	1100 1010	1011 0101
2	Ignore the first 9 bits and copy the remaining 23 bits	X	110 0000	1100 1010	1011 0101
3	First bit X = 0 for Internet; X = 1 for other	0	110 0000	1100 1010	1011 0101
4	Written in hexadecimal		60	CA	B5
5	MAC address in hexadecimal	01 : 00 : 5E : E0 : CA : B5			
6	Drop last 24 bits	01 : 00 : 5E :			
7	Copy the multicast bits	01 : 00 : 5E : 60 : CA : B5			
8	MAC frame destination address 01:00:5E:60:CA:B5 corresponds to multicast IPv4 address 232.224.202.181				

Note that this process means that there are 32 (2^5) IPv4 multicast addresses that could map to the same MAC multicast addresses. For example, multicast IPv4 addresses 224.8.7.6 and 229.136.7.6 translate to the same MAC address (0x01-00-5E-08-07-06). This is a real concern, and because the host could be interested in frames sent to both of those multicast groups, the IP software must reject one or the other.



NOTE: This “collision” problem does not exist in IPv6 because of the way IPv6 handles multicast groups, but it is always a concern in IPv4. The procedure for placing IPv6 multicast packets inside multicast frames is nearly identical to that for IPv4, except for the MAC destination address 0x3333 prefix (and the lack of “collisions”).

Once the MAC address for the multicast group is determined, the host's operating system essentially orders the LAN interface card to join or leave the multicast group. Once joined to a multicast group, the host accepts frames sent to the multicast address as well as the host's unicast address and ignores other multicast group's frames. It is possible for a host to join and receive multicast content from more than one group at the same time, of course.

Multicast Interface Lists

To avoid multicast routing loops, every multicast routing device must always be aware of the interface that leads to the source of that multicast group content by the shortest path. This is the upstream (incoming) interface, and packets are never to be forwarded back toward a multicast source. All other interfaces are potential downstream (outgoing) interfaces, depending on the number of branches on the distribution tree.

routing devices closely monitor the status of the incoming and outgoing interfaces, a process that determines the *multicast forwarding state*. A routing device with a multicast forwarding state for a particular multicast group is essentially “turned on” for that group's

content. Interfaces on the routing device's outgoing interface list send copies of the group's packets received on the incoming interface list for that group. The incoming and outgoing interface lists might be different for different multicast groups.

The multicast forwarding state in a routing device is usually written in either (S,G) or (*,G) notation. These are pronounced “ess comma gee” and “star comma gee,” respectively. In (S,G), the S refers to the unicast IP address of the source for the multicast traffic, and the G refers to the particular multicast group IP address for which S is the source. All multicast packets sent from this source have S as the source address and G as the destination address.

The asterisk (*) in the (*,G) notation is a wildcard indicating that the state applies to any multicast application source sending to group G. So, if two sources are originating exactly the same content for multicast group 224.1.1.2, a routing device could use (*,224.1.1.2) to represent the state of a routing device forwarding traffic from both sources to the group.

Multicast Routing Protocols

Multicast routing protocols enable a collection of multicast routing devices to build (join) distribution trees when a host on a directly attached subnet, typically a LAN, wants to receive traffic from a certain multicast group, prune branches, locate sources and groups, and prevent routing loops.

There are several multicast routing protocols:

- **Distance Vector Multicast Routing Protocol (DVMRP)**—The first of the multicast routing protocols and hampered by a number of limitations that make this method unattractive for large-scale Internet use. DVMRP is a dense-mode-only protocol, and uses the flood-and-prune or implicit join method to deliver traffic everywhere and then determine where the uninterested receivers are. DVMRP uses source-based distribution trees in the form (S,G), and builds its own multicast routing tables for RPF checks.
- **Multicast OSPF (MOSPF)**—Extends OSPF for multicast use, but only for dense mode. However, MOSPF has an explicit join message, so routing devices do not have to flood their entire domain with multicast traffic from every source. MOSPF uses source-based distribution trees in the form (S,G).
- ***Bidirectional PIM mode***—A variation of PIM. Bidirectional PIM builds bidirectional shared trees that are rooted at a rendezvous point (RP) address. Bidirectional traffic does not switch to shortest path trees as in PIM-SM and is therefore optimized for routing state size instead of path length. This means that the end-to-end latency might be longer compared to PIM sparse mode. Bidirectional PIM routes are always wildcard-source (*,G) routes. The protocol eliminates the need for (S,G) routes and data-triggered events. The bidirectional (*,G) group trees carry traffic both upstream from senders toward the RP, and downstream from the RP to receivers. As a consequence, the strict reverse path forwarding (RPF)-based rules found in other PIM modes do not apply to bidirectional PIM. Instead, bidirectional PIM (*,G) routes forward traffic from all sources and the RP. Bidirectional PIM routing devices must have the ability to accept traffic on many potential incoming interfaces. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Bidirectional PIM is recommended in deployments with many dispersed sources and many dispersed receivers.

- *PIM dense mode*—In this mode of PIM, the assumption is that almost all possible subnets have at least one receiver wanting to receive the multicast traffic from a source, so the network is *flooded* with traffic on all possible branches, then pruned back when branches do not express an interest in receiving the packets, explicitly (by message) or implicitly (time-out silence). This is the *dense mode* of multicast operation. LANs are appropriate networks for dense-mode operation. Some multicast routing protocols, especially older ones, support only dense-mode operation, which makes them inappropriate for use on the Internet. In contrast to DVMRP and MOSPF, PIM dense mode allows a routing device to use any unicast routing protocol and performs RPF checks using the unicast routing table. PIM dense mode has an implicit join message, so routing devices use the flood-and-prune method to deliver traffic everywhere and then determine where the uninterested receivers are. PIM dense mode uses source-based distribution trees in the form (S,G), as do all dense-mode protocols. PIM also supports sparse-dense mode, with mixed sparse and dense groups, but there is no special notation for that operational mode. If *sparse-dense mode* is supported, the multicast routing protocol allows some multicast groups to be sparse and other groups to be dense.
- *PIM sparse mode*—In this mode of PIM, the assumption is that very few of the possible receivers want packets from each source, so the network establishes and sends packets only on branches that have at least one leaf indicating (by message) an interest in the traffic. This multicast protocol allows a routing device to use any unicast routing protocol and performs reverse-path forwarding (RPF) checks using the unicast routing table. PIM sparse mode has an *explicit* join message, so routing devices determine where the interested receivers are and send join messages upstream to their neighbors, building trees from receivers to the rendezvous point (RP). PIM sparse mode uses an RP routing device as the initial source of multicast group traffic and therefore builds distribution trees in the form (*,G), as do all sparse-mode protocols. PIM sparse mode migrates to an (S,G) source-based tree if that path is shorter than through the RP for a particular multicast group's traffic. WANs are appropriate networks for sparse-mode operation, and indeed a common multicast guideline is not to run dense mode on a WAN under any circumstances.
- *Core Based Trees (CBT)*—Shares all of the characteristics of PIM sparse mode (sparse mode, explicit join, and shared (*,G) trees), but is said to be more efficient at finding sources than PIM sparse mode. CBT is rarely encountered outside academic discussions. There are no large-scale deployments of CBT, commercial or otherwise.
- *PIM source-specific multicast (SSM)*—Enhancement to PIM sparse mode that allows a client to receive multicast traffic directly from the source, without the help of an RP. Used with IGMPv3 to create a shortest-path tree between receiver and source.
- *IGMPv1*—The original protocol defined in RFC 1112, *Host Extensions for IP Multicasting*. IGMPv1 sends an explicit join message to the routing device, but uses a timeout to determine when hosts leave a group. Three versions of the Internet Group Management Protocol (IGMP) run between receiver hosts and routing devices.
- *IGMPv2*—Defined in RFC 2236, *Internet Group Management Protocol, Version 2*. Among other features, IGMPv2 adds an explicit leave message to the join message.
- *IGMPv3*—Defined in RFC 3376, *Internet Group Management Protocol, Version 3*. Among other features, IGMPv3 optimizes support for a single source of content for a multicast

group, or source-specific multicast (SSM). Used with PIM SSM to create a shortest-path tree between receiver and source.

- Bootstrap Router (BSR) and Auto-Rendezvous Point (RP)—Allow sparse-mode routing protocols to find RPs within the routing domain (autonomous system, or AS). RP addresses can also be statically configured.
- Multicast Source Discovery Protocol (MSDP)—Allows groups located in one multicast routing domain to find RPs in other routing domains. MSDP is not used on an RP if all receivers and sources are located in the same routing domain. Typically runs on the same routing device as PIM sparse mode RP. Not appropriate if all receivers and sources are located in the same routing domain.
- Session Announcement Protocol (SAP) and Session Description Protocol (SDP)—Display multicast session names and correlate the names with multicast traffic. SDP is a session directory protocol that advertises multimedia conference sessions and communicates setup information to participants who want to join the session. A client commonly uses SDP to announce a conference session by periodically multicasting an announcement packet to a well-known multicast address and port using SAP.
- Pragmatic General Multicast (PGM)—Special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic. PGM allows a receiver to detect missing information in all cases and request replacement information if the receiver application requires it.

The differences among the multicast routing protocols are summarized in [Table 3 on page 34](#).

Table 3: Multicast Routing Protocols Compared

Multicast Routing Protocol	Dense Mode	Sparse Mode	Implicit Join	Explicit Join	(S,G) SBT	(*G) Shared Tree
DVMRP	Yes	No	Yes	No	Yes	No
MOSPF	Yes	No	No	Yes	Yes	No
PIM dense mode	Yes	No	Yes	No	Yes	No
PIM sparse mode	No	Yes	No	Yes	Yes, maybe	Yes, initially
Bidirectional PIM	No	No	No	Yes	No	Yes
CBT	No	Yes	No	Yes	No	Yes
SSM	No	Yes	No	Yes	Yes, maybe	Yes, initially
IGMPv1	No	Yes	No	Yes	Yes, maybe	Yes, initially
IGMPv2	No	Yes	No	Yes	Yes, maybe	Yes, initially
IGMPv3	No	Yes	No	Yes	Yes, maybe	Yes, initially

Table 3: Multicast Routing Protocols Compared (*continued*)

Multicast Routing Protocol	Dense Mode	Sparse Mode	Implicit Join	Explicit Join	(S,G) SBT	(*G) Shared Tree
BSR and Auto-RP	No	Yes	No	Yes	Yes, maybe	Yes, initially
MSDP	No	Yes	No	Yes	Yes, maybe	Yes, initially

It is important to realize that retransmissions due to a high bit-error rate on a link or overloaded routing device can make multicast as inefficient as repeated unicast. Therefore, there is a trade-off in many multicast applications regarding the session support provided by the Transmission Control Protocol (TCP) (but TCP always resends missing segments), or the simple drop-and-continue strategy of the User Datagram Protocol (UDP) datagram service (but reordering can become an issue). Modern multicast uses UDP almost exclusively.

T Series Router Multicast Performance

The Juniper Networks T Series Core Routers handle extreme multicast packet replication requirements with a minimum of router load. Each memory component replicates a multicast packet twice at most. Even in the worst-case scenario involving maximum fan-out, when 1 input port and 63 output ports need a copy of the packet, the T Series routing platform copies a multicast packet only six times. Most multicast distribution trees are much sparser, so in many cases only two or three replications are necessary. In no case does the T Series architecture have an impact on multicast performance, even with the largest multicast fan-out requirements.

PART 1

Managing Group Membership

- [Using IGMP on page 39](#)
- [Using IGMP Snooping on page 65](#)
- [Using MLD on page 75](#)
- [Using MLD Snooping on page 101](#)

CHAPTER 2

Using IGMP

- [Understanding Group Membership Protocols on page 39](#)
- [Understanding IGMP on page 40](#)
- [Configuring IGMP on page 42](#)
- [Enabling IGMP on page 43](#)
- [Changing the IGMP Version on page 44](#)
- [Modifying the IGMP Host-Query Message Interval on page 45](#)
- [Modifying the IGMP Last-Member Query Interval on page 46](#)
- [Specifying Immediate-Leave Host Removal for IGMP on page 47](#)
- [Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 48](#)
- [Accepting IGMP Messages from Remote Subnetworks on page 49](#)
- [Modifying the IGMP Query Response Interval on page 50](#)
- [Modifying the IGMP Robustness Variable on page 51](#)
- [Limiting the Maximum IGMP Message Rate on page 52](#)
- [Enabling IGMP Static Group Membership on page 52](#)
- [Recording IGMP Join and Leave Events on page 59](#)
- [Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 60](#)
- [Tracing IGMP Protocol Traffic on page 61](#)
- [Disabling IGMP on page 63](#)

Understanding Group Membership Protocols

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

and distribution tree from source to leaf. So, different multicast protocols are used for the host-router portion and the router-router portion of the multicast network.

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

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

- IGMPv1—The original protocol defined in RFC 1112. An explicit join message is sent to the routing device, but a timeout is used to determine when hosts leave a group. This process wastes processing cycles on the routing device, especially on older or smaller routing devices.
- IGMPv2—Defined in RFC 2236. Among other features, IGMPv2 adds an explicit leave message to the join message so that routing devices can more easily determine when a group has no interested listeners on a LAN.
- IGMPv3—Defined in RFC 3376. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or *source-specific multicast (SSM)*.
- MLDv1—Defined in RFC 2710. MLDv1 is similar to IGMPv2.
- MLDv2—Defined in RFC 3810. MLDv2 similar to IGMPv3.

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

**Related
Documentation**

- [Examples: Configuring MLD on page 78](#)

Understanding IGMP

The Internet Group Management Protocol (IGMP) manages the membership of hosts and routing devices in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routing devices. Multicast routing devices use IGMP to learn, for each of their attached physical networks, which groups have members.

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

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

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

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

A route entry can include such fields as:

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

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

The outgoing interfaces point to the neighboring downstream routing devices that have sent join and prune messages toward the RP as well as the directly connected hosts that have requested membership to group G.

5. This state creates a shared, RP-centered, distribution tree that reaches all group members.

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

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

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

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

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

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

**Related
Documentation**

- *Supported IP Multicast Protocol Standards*
- *Configuring IGMP*

Configuring IGMP

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements. See *Configuring the Session Announcement Protocol*.

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

```
igmp {  
  accounting;  
  interface interface-name {  
    disable;  
    (accounting | no-accounting);  
    group-policy [ policy-names ];  
    immediate-leave;  
    oif-map map-name;
```

```

promiscuous-mode;
ssm-map ssm-map-name;
static {
    group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
            source-count number;
            source-increment increment;
        }
    }
}
version version;
}
query-interval seconds;
query-last-member-interval seconds;
query-response-interval seconds;
robust-count number;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}

```

You can include this statement at the following hierarchy levels:

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

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



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

Enabling IGMP

The Internet Group Management Protocol (IGMP) manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use IGMP to learn which groups have members on each of their attached physical

networks. IGMP must be enabled for the router to receive IPv4 multicast packets. IGMP is only needed for IPv4 networks, because multicast is handled differently in IPv6 networks. IGMP is automatically enabled on all IPv4 interfaces on which you configure PIM and on all IPv4 broadcast interfaces when you configure DVMRP.

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

To explicitly enable IGMP:

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

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

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

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

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

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

4. Verify the configuration.

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

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

- Related Documentation**
- [Understanding IGMP on page 40](#)
 - [Disabling IGMP on page 63](#)
 - [show igmp interface on page 435](#)

Changing the IGMP Version

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

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

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

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

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

To change to IGMPv3 for SSM functionality:

1. Configure the IGMP interface.

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

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



CAUTION: On MX Series platforms, IGMPv2 and IGMPv3 cannot be configured together on the same interface. Configuring both together causes unexpected behavior in multicast traffic forwarding.

Related Documentation

- [Understanding IGMP on page 40](#)
- [show pim interfaces on page 517](#)
- [show igmp statistics on page 439](#)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 3376, *Internet Group Management Protocol, Version 3*

Modifying the IGMP Host-Query Message Interval

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

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

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

To modify the query interval:

1. Configure the interval.

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

The value can be from 1 through 1024 seconds.

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

Related Documentation

- [Understanding IGMP on page 40](#)
- [Modifying the IGMP Query Response Interval on page 50](#)
- [Modifying the IGMP Robustness Variable on page 51](#)
- [show igmp interface on page 435](#)
- [show igmp statistics on page 439](#)

Modifying the IGMP Last-Member Query Interval

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

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

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

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

To modify this interval:

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

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

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



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

**Related
Documentation**

- [Modifying the IGMP Robustness Variable on page 51](#)
- [show pim interfaces on page 517](#)

Specifying Immediate-Leave Host Removal for IGMP

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

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

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

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



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

To enable immediate leave on an interface:

1. Configure immediate leave on the IGMP interface.

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

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

Related Documentation

- [Understanding IGMP on page 40](#)
- [show igmp interface on page 435](#)

Filtering Unwanted IGMP Reports at the IGMP Interface Level

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

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



CAUTION: On MX Series platforms, IGMPv2 and IGMPv3 cannot be configured together on the same interface. Configuring both together causes unexpected behavior in multicast traffic forwarding.

To filter unwanted IGMP reports:

1. Configure an IGMPv2 policy.

```
[edit policy-statement reject_policy_v2]
```

```

user@host# set from route-filter 224.1.1.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set then reject

```

2. Configure an IGMPv3 policy.

```

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

```

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

```

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

```

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

Related Documentation

- [Understanding IGMP on page 40](#)
- [Example: Configuring Policy Chains and Route Filters](#)
- [show igmp statistics on page 439](#)

Accepting IGMP Messages from Remote Subnetworks

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



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



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

To enable IGMP promiscuous mode on an interface:

1. Configure the IGMP interface.

```

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

```

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

**Related
Documentation**

- [Understanding IGMP on page 40](#)
- *Configuring the Loopback Interface* in the *Junos OS Network Interfaces Library for Routing Devices*
- [show igmp interface on page 435](#)
- [show igmp statistics on page 439](#)

Modifying the IGMP Query Response Interval

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

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

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

To modify the query response interval:

1. Configure the interval.

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

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

**Related
Documentation**

- [Understanding IGMP on page 40](#)
- [Modifying the IGMP Host-Query Message Interval on page 45](#)

- [Modifying the IGMP Robustness Variable on page 51](#)
- [show igmp interface on page 435](#)
- [show igmp statistics on page 439](#)

Modifying the IGMP Robustness Variable

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

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

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

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

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

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

The number can be from 2 through 10.

To change the value of the robustness variable:

1. Configure the robust count.

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

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

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

Related Documentation

- [Modifying the IGMP Host-Query Message Interval on page 45](#)
- [Modifying the IGMP Query Response Interval on page 50](#)
- [Modifying the IGMP Last-Member Query Interval on page 46](#)
- [show pim interfaces on page 517](#)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 3376, *Internet Group Management Protocol, Version 3*

Limiting the Maximum IGMP Message Rate

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

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

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

Related Documentation

- [maximum-transmit-rate \(Protocols IGMP\) on page 262](#)

Enabling IGMP Static Group Membership

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

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

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

In this example, you create static group 225.1.1.1.

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


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

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

```
user@host> show configuration protocol igmp

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

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

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



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

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

In this example, you create three groups.

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

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

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

```
user@host> show configuration protocol igmp

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

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

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

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

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

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

```
[edit protocols igmp]
```

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

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

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      group-increment 0.0.0.2;
      group-count 3;
    }
  }
}
```

- After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.3, and 225.1.1.5 have been created.

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

```

Interface: fe-0/1/2
  Group: 225.1.1.1
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 225.1.1.3
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 225.1.1.5
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static

```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, and your network is operating in source-specific multicast (SSM) mode, you can also specify that the multicast source address be accepted. This is useful when you want to test forwarding to multicast receivers from a specific multicast source.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you create group 225.1.1.1 and accept IP address 10.0.0.2 as the only source.

1. On the DR, configure the source address by including the **source** statement and specifying the IPv4 address of the source host.

```

[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2

```

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

```

user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      source 10.0.0.2;
    }
  }
}

```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that source 10.0.0.2 has been accepted.

```

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

```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of multicast sources be automatically accepted. This is useful when you want to test forwarding to multicast receivers from more than one specified multicast source.

In this example, you create group 255.1.1.1 and accept addresses 10.0.0.2, 10.0.0.3, and 10.0.0.4 as the sources.

1. On the DR, configure the number of multicast source addresses to be accepted by including the **source-count** statement and specifying the number of sources to be accepted.

```
[edit protocols igmp]
```

```
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count 3
```

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

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {  
  version 3;  
  static {  
    group 225.1.1.1 {  
      source 10.0.0.2 {  
        source-count 3;  
      }  
    }  
  }  
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.3, and 10.0.0.4 have been accepted.

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

```
Interface: fe-0/1/2  
  Group: 225.1.1.1  
    Source: 10.0.0.2  
    Last reported by: Local  
    Timeout: 0 Type: Static  
  Group: 225.1.1.1  
    Source: 10.0.0.3  
    Last reported by: Local  
    Timeout: 0 Type: Static  
  Group: 225.1.1.1  
    Source: 10.0.0.4  
    Last reported by: Local  
    Timeout: 0 Type: Static
```

When you configure static groups on an interface on which you want to receive multicast traffic, and specify that a number of multicast sources be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and you do not want the source addresses to be sequential.

In this example, you create group 225.1.1.1 and accept addresses 10.0.0.2, 10.0.0.4, and 10.0.0.6 as the sources.

1. Configure the multicast source address increment by including the **source-increment** statement and specifying the number by which the address should be incremented for each source. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
```

```
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count 3 source-increment 0.0.0.2
```

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

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      source 10.0.0.2 {
        source-count 3;
        source-increment 0.0.0.2;
      }
    }
  }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.4, and 10.0.0.6 have been accepted.

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

```
Interface: fe-0/1/2
  Group: 225.1.1.1
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 225.1.1.1
    Source: 10.0.0.4
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 225.1.1.1
    Source: 10.0.0.6
    Last reported by: Local
    Timeout: 0 Type: Static
```

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the source address configured. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the source address configured.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you exclude address 10.0.0.2 as a source for group 225.1.1.1.

1. On the DR, configure a multicast static group to operate in exclude mode by including the **exclude** statement and specifying which IPv4 source address to exclude.

```
[edit protocols igmp]
```

```
user@host# set interface fe-0/1/2 static group 225.1.1.1 exclude source 10.0.0.2
```

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

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      exclude;
      source 10.0.0.2;
    }
  }
}
```

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

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

```
Interface: fe-0/1/2
Group: 225.1.1.1
Group mode: Exclude
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
```

Related Documentation

- [Enabling MLD Static Group Membership on page 89](#)
- [group \(Protocols IGMP\) on page 255](#)
- [group-count \(Protocols IGMP\) on page 256](#)
- [group-increment \(Protocols IGMP\) on page 256](#)
- [source-count \(Protocols IGMP\) on page 270](#)

- [source-increment \(Protocols IGMP\) on page 271](#)
- [static \(Protocols IGMP\) on page 272](#)

Recording IGMP Join and Leave Events

To determine whether IGMP tuning is needed in a network, you can configure the routing device to record IGMP join and leave events. You can record events globally for the routing device or for individual interfaces.

[Table 4 on page 59](#) describes the recordable IGMP events.

Table 4: IGMP Event Messages

ERRMSG Tag	Definition
RPD_IGMP_JOIN	Records IGMP join events.
RPD_IGMP_LEAVE	Records IGMP leave events.
RPD_IGMP_ACCOUNTING_ON	Records when IGMP accounting is enabled on an IGMP interface.
RPD_IGMP_ACCOUNTING_OFF	Records when IGMP accounting is disabled on an IGMP interface.
RPD_IGMP_MEMBERSHIP_TIMEOUT	Records IGMP membership timeout events.

To enable IGMP accounting:

1. Enable accounting globally or on an IGMP interface. This example shows both options.

```
[edit protocols igmp]
user@host# set accounting
user@host# set interface fe-0/1/0.2 accounting
```

2. Configure the events to be recorded and filter the events to a system log file with a descriptive filename, such as `igmp-events`.

```
[edit system syslog file igmp-events]
user@host# set any info
user@host# set match ".*RPD_IGMP_JOIN.* | .*RPD_IGMP_LEAVE.* |
.*RPD_IGMP_ACCOUNTING.* | .*RPD_IGMP_MEMBERSHIP_TIMEOUT.*"
```

3. Periodically archive the log file.

This example rotates the file size when it reaches 100 KB and keeps three files.

```
[edit system syslog file igmp-events]
user@host# set archive size 100000
user@host# set archive files 3
user@host# set archive archive-sites "ftp://user@host1//var/tmp" password
"anonymous"
user@host# set archive archive-sites "ftp://user@host2//var/tmp" password "test"
user@host# set archive transfer-interval 24
user@host# set archive start-time 2011-01-07:12:30
```

4. You can monitor the system log file as entries are added to the file by running the **monitor start** and **monitor stop** commands.

```
user@host> monitor start igmp-events
```

```
*** igmp-events ***
```

```
Apr 16 13:08:23 host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command  
'run monitor start igmp-events '  
monitor
```

**Related
Documentation**

- [Understanding IGMP on page 40](#)
- [Specifying Log File Size, Number, and Archiving Properties](#)

Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

The **group-limit** statement enables you to limit the number of IGMP multicast group joins for logical interfaces. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for IGMP multicast groups, keep the following in mind:

- Each any-source group (*G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in IGMPv3 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The groups must then request to rejoin the network (up to the newly configured group limit).
- You can dynamically limit multicast groups on IGMP logical interfaces using dynamic profiles.

Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for IGMP multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of IGMP multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The **group-threshold** statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a

percentage of the group limit, so you must configure the **group-limit** statement to configure a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs the warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the **log-interval** statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for IGMP multicast group joins.

To limit multicast group joins on an IGMP logical interface:

1. Access the logical interface at the IGMP protocol hierarchy level.

```
[edit]
user@host# edit protocols igmp interface interface-name
```

2. Specify the group limit for the interface.

```
[edit protocols igmp interface interface-name]
user@host# set group-limit limit
```

3. (Optional) Configure the threshold at which a warning message is logged.

```
[edit protocols igmp interface interface-name]
user@host# set group-threshold value
```

4. (Optional) Configure the amount of time between log messages.

```
[edit protocols igmp interface interface-name]
user@host# set log-interval seconds
```

To confirm your configuration, use the **show protocols igmp** command. To verify the operation of IGMP on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the **show igmp interface** command.

Related Documentation

- [Enabling IGMP Static Group Membership on page 52](#)

Tracing IGMP Protocol Traffic

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

Flag	Description
all	Trace all operations.
client-notification	Trace notifications.

Flag	Description
general	Trace general flow.
group	Trace group operations.
host-notification	Trace host notifications.
leave	Trace leave group messages (IGMPv2 only).
mtrace	Trace mtrace packets. Use the mtrace command to troubleshoot the software.
normal	Trace normal events.
packets	Trace all IGMP packets.
policy	Trace policy processing.
query	Trace IGMP membership query messages, including general and group-specific queries.
report	Trace membership report messages.
route	Trace routing information.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

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

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

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

2. Configure the filename for the IGMP trace file.

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

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

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

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

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

5. (Optional) Enable unrestricted file access.

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

6. Configure tracing flags. Suppose you are troubleshooting issues with a particular multicast group. The following example shows how to flag all events for packets associated with the group IP address.

```
[edit protocols igmp traceoptions]  
user@host# set flag group | match 232.1.1.2
```

7. View the trace file.

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

- Related Documentation**
- [Understanding IGMP on page 40](#)
 - *Tracing and Logging Junos OS Operations*
 - [mtrace on page 473](#)

Disabling IGMP

To disable IGMP on an interface, include the **disable** statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

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

- Related Documentation**
- [Enabling IGMP on page 43](#)

CHAPTER 3

Using IGMP Snooping

- [IGMP Snooping Overview on page 65](#)
- [Configuring IGMP Snooping on page 68](#)
- [Configuring VLAN-Specific IGMP Snooping Parameters on page 69](#)
- [Example: Configuring IGMP Snooping on page 70](#)
- [Monitoring IGMP Snooping on page 72](#)
- [Verifying the IGMP Snooping Group Timeout Value on page 73](#)

IGMP Snooping Overview

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

This IGMP snooping topic includes:

- [How IGMP Snooping Works on page 65](#)
- [How IGMP Snooping Works with Routed VLAN Interfaces on page 66](#)
- [How Hosts Join and Leave Multicast Groups on page 66](#)
- [IGMP Snooping and Forwarding Interfaces on page 67](#)
- [General Forwarding Rules on page 67](#)

How IGMP Snooping Works

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

IGMP snooping prevents this flooding. When you enable IGMP snooping, the switch monitors IGMP packets between receivers and multicast routers and uses the content

of the packets to build a multicast cache table—a database of multicast groups and the interfaces that are connected to members of the groups. When the switch receives multicast packets, it uses the cache table to selectively forward the traffic to only the interfaces that are connected to members of the appropriate multicast groups.



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



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

How IGMP Snooping Works with Routed VLAN Interfaces

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

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

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

How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in two ways:

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

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

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

IGMP Snooping and Forwarding Interfaces

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

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

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

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



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

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

General Forwarding Rules

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

IGMP traffic is forwarded as follows:

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

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

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

**Related
Documentation**

- [Example: Configuring IGMP Snooping on page 70](#)
- [Configuring IGMP Snooping on page 68](#)
- [Monitoring IGMP Snooping on page 72](#)
- [Configuring IGMP on page 42](#)
- RFC 3171, *IANA Guidelines for IPv4 Multicast Address Assignments*
- IGMPv1—See RFC 1112, *Host extensions for IP multicasting*.
- IGMPv2—See RFC 2236, *Internet Group Management Protocol, Version 2*.
- IGMPv3—See RFC 3376, *Internet Group Management Protocol, Version 3*.

Configuring IGMP Snooping

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



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

To enable IGMP snooping and configure individual options as needed for your network by using the CLI:

1. Enable IGMP snooping on a VLAN:

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

2. Configure the switch to immediately remove group membership from interfaces on a VLAN when it receives a leave message through that VLAN, and have it not forward any membership queries for the multicast group to the VLAN (IGMPv2 only):

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

3. Configure an interface to belong to a multicast group:


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

4. Configure an interface to forward IGMP queries received from multicast routers.

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

5. Configure the switch to wait for four timeout intervals before timing out a multicast group on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name robust-count 4
```

6. If you want a standalone switch to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name l2-querier source-address source address
```

The switch uses the address that you configure as the source address in the IGMP queries that it sends. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is greater (a higher number) than the IP addresses for those routers on the network. This ensures that switch is always the IGMP querier on the network.

- Related Documentation**
- [IGMP Snooping Overview on page 65](#)
 - [Example: Configuring IGMP Snooping on page 70](#)
 - [Monitoring IGMP Snooping on page 72](#)

Configuring VLAN-Specific IGMP Snooping Parameters

All of the IGMP snooping statements configured with the **igmp-snooping** statement, with the exception of the **traceoptions** statement, can be qualified with the same statement at the VLAN level. To configure IGMP snooping parameters at the VLAN level, include the **vlan** statement:

```
vlan vlan-id;
  immediate-leave;
  interface interface-name {
    group-limit limit;
    host-only-interface;
    multicast-router-interface;
    static {
      group ip-address {
        source ip-address;
      }
    }
  }
  proxy {
    source-address ip-address;
  }
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
```

```
}
```

You can include this statement at the following hierarchy levels:

- [edit bridge-domains *bridge-domain-name* protocols igmp-snooping]
- [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping]

Related Documentation

- [Layer 2 Frames and IPv4 Multicast Addresses on page 29](#)
- [Understanding Multicast Snooping](#)

Example: Configuring IGMP Snooping

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

This example describes how to configure IGMP snooping:

- [Requirements on page 70](#)
- [Overview and Topology on page 70](#)
- [Configuration on page 71](#)

Requirements

This example requires Junos OS Release 11.1 or later on a QFX Series product.

Before you configure IGMP snooping, be sure you have:

- Configured the **employee-vlan** VLAN
- Assigned interfaces **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/3** to **employee-vlan**

Overview and Topology

In this example you configure an interface to receive multicast traffic from a source and configure some multicast-related behavior for downstream interfaces. The example assumes that IGMP snooping was previously disabled for the VLAN.

[Table 5 on page 70](#) shows the components of the topology for this example.

Table 5: Components of the IGMP Snooping Topology

Components	Settings
VLAN name	employee-vlan , tag 20
Interfaces in employee-vlan	ge-0/0/1 , ge-0/0/2 , ge-0/0/3

Table 5: Components of the IGMP Snooping Topology (*continued*)

Components	Settings
Multicast IP address for employee-vlan	225.100.100.100

Configuration

To configure basic IGMP snooping on a switch:

CLI Quick Configuration

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

```
[edit protocols]
set igmp-snooping vlan employee-vlan
set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group 225.100.100.100
set igmp-snooping vlan employee-vlan interface ge-0/0/2 multicast-router-interface
set igmp-snooping vlan employee-vlan robust-count 4
```

Step-by-Step Procedure

Configure IGMP snooping:

1. Enable and configure IGMP snooping on the VLAN **employee-vlan**:

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

2. Configure an interface to belong to a multicast group:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group
225.100.100.100
```

3. Configure an interface to forward IGMP queries received from multicast routers.

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/2
multicast-router-interface
```

4. Configure the switch to wait for four timeout intervals before timing out a multicast group on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan robust-count 4
```

Results

Check the results of the configuration:

```
user@switch# show protocols igmp-snooping
vlan employee-vlan {
  robust-count 4;
}
interface ge-0/0/2 {
  multicast-router-interface;
}
interface ge-0/0/3 {
  static {
    group 225.100.100.100;
  }
}
```

- Related Documentation**
- [IGMP Snooping Overview on page 65](#)
 - [Configuring IGMP Snooping on page 68](#)
 - [Changing the IGMP Snooping Group Timeout Value](#)
 - [Monitoring IGMP Snooping on page 72](#)
 - [Example: Setting Up Bridging with Multiple VLANs.](#)

Monitoring IGMP Snooping

Purpose Use the monitoring feature to view status and information about the IGMP snooping configuration.

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

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

Meaning [Table 6 on page 72](#) summarizes the IGMP snooping details displayed.

Table 6: Summary of IGMP Snooping Output Fields

Field	Values
IGMP Snooping Monitor	
VLAN	VLAN for which IGMP snooping is enabled.
Interfaces	Interface connected to a multicast router.
Groups	Number of the multicast groups learned by the VLAN.
MRouters	Multicast router.
Receivers	Multicast receiver.
IGMP Route Information	
VLAN	VLAN for which IGMP snooping is enabled.
Next-Hop	Next hop assigned by the switch after performing the route lookup.
Group	Multicast groups learned by the VLAN.

- Related Documentation**
- [IGMP Snooping Overview on page 65](#)

- [Example: Configuring IGMP Snooping on page 70](#)
- [Configuring IGMP Snooping on page 68](#)
- [Changing the IGMP Snooping Group Timeout Value](#)

Verifying the IGMP Snooping Group Timeout Value

Purpose Verify that the IGMP snooping group timeout value has been changed correctly from its default value.

Action Display the IGMP snooping membership information, which contains the group timeout value that was derived from the IGMP configuration:

```
user@switch> show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
  Group: 225.0.0.1
    Receiver count: 1, Flags: <v2-hosts>
      ge-0/0/15.0 Uptime: 00:00:05 timeout: 510
```

Meaning The IGMP snooping group timeout value determines how long a switch waits to receive an IGMP query from a multicast router before removing a multicast group from its multicast cache table. When you enable IGMP snooping, the default IGMP snooping group timeout value of 260 seconds is applied to all VLANs, which means that the switch waits 260 seconds to receive an IGMP query before removing a multicast group from its multicast cache table. You can change the timeout value by using the **robust-count** option.

Related Documentation

- [Changing the IGMP Snooping Group Timeout Value](#)

CHAPTER 4

Using MLD

- [Understanding MLD on page 75](#)
- [Examples: Configuring MLD on page 78](#)

Understanding MLD

The Multicast Listener Discovery (MLD) Protocol manages the membership of hosts and routers in multicast groups. IP version 6 (IPv6) multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol.


MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

In include mode, the receiver specifies the source or sources it is interested in receiving the multicast group traffic from. Exclude mode works the opposite of include mode. It allows the receiver to specify the source or sources it is not interested in receiving the multicast group traffic from.

For each attached network, a multicast routing device can be either a querier or a nonquerier. A querier routing device, usually one per subnet, solicits group membership information by transmitting MLD queries. When a host reports to the querier routing device that it has interested listeners, the querier routing device forwards the membership information to the rendezvous point (RP) routing device by means of the receiver's (host's) designated router (DR). This builds the rendezvous-point tree (RPT) connecting the host with interested listeners to the RP routing device. The RPT is the initial path used by the sender to transmit information to the interested listeners. Nonquerier routing devices do not transmit MLD queries on a subnet but can do so if the querier routing device fails.

All MLD-configured routing devices start as querier routing devices on each attached subnet (see [Figure 3 on page 76](#)). The querier routing device on the right is the receiver's DR.

Figure 3: Routing Devices Start Up on a Subnet




To elect the querier routing device, the routing devices exchange query messages containing their IPv6 source addresses. If a routing device hears a query message whose IPv6 source address is numerically lower than its own selected address, it becomes a nonquerier. In [Figure 4 on page 76](#), the routing device on the left has a source address numerically lower than the one on the right and therefore becomes the querier routing device.




NOTE: In the practical application of MLD, several routing devices on a subnet are nonqueriers. If the elected querier routing device fails, query messages are exchanged among the remaining routing devices. The routing device with the lowest IPv6 source address becomes the new querier routing device. The IPv6 Neighbor Discovery Protocol (NDP) implementation drops incoming Neighbor Announcement (NA) messages that have a broadcast or multicast address in the target link-layer address option. This behavior is recommended by RFC 2461.

Figure 4: Querier Routing Device Is Determined




The querier routing device sends general MLD queries on the **link-scope all-nodes** multicast address FF02::1 at short intervals to all attached subnets to solicit group membership information (see [Figure 5 on page 77](#)). Within the query message is the *maximum response delay* value, specifying the maximum allowed delay for the host to respond with a report message.

Figure 5: General Query Message Is Issued




If interested listeners are attached to the host receiving the query, the host sends a report containing the host's IPv6 address to the routing device (see [Figure 6 on page 77](#)). If the reported address is not yet in the routing device's list of multicast addresses with interested listeners, the address is added to the list and a timer is set for the address. If the address is already on the list, the timer is reset. The host's address is transmitted to the RP in the PIM domain.

Figure 6: Reports Are Received by the Querier Routing Device




If the host has no interested multicast listeners, it sends a done message to the querier routing device. On receipt, the querier routing device issues a multicast address-specific query containing the last **listener query interval** value to the multicast address of the host. If the routing device does not receive a report from the multicast address, it removes the multicast address from the list and notifies the RP in the PIM domain of its removal (see [Figure 7 on page 77](#)).

Figure 7: Host Has No Interested Receivers and Sends a Done Message to Routing Device



If a done message is not received by the querier routing device, the querier routing device continues to send multicast address-specific queries. If the timer set for the address on receipt of the last report expires, the querier routing device assumes there are no longer interested listeners on that subnet, removes the multicast address from the list, and notifies the RP in the PIM domain of its removal (see [Figure 8 on page 78](#)).

Figure 8: Host Address Timer Expires and Address Is Removed from Multicast Address List



Related Documentation

- [Enabling MLD on page 82](#)
- [Example: Recording MLD Join and Leave Events on page 96](#)
- [Example: Modifying the MLD Robustness Variable on page 87](#)

Examples: Configuring MLD

- [Understanding MLD on page 78](#)
- [Configuring MLD on page 81](#)
- [Enabling MLD on page 82](#)
- [Modifying the MLD Version on page 83](#)
- [Modifying the MLD Host-Query Message Interval on page 83](#)
- [Modifying the MLD Query Response Interval on page 84](#)
- [Modifying the MLD Last-Member Query Interval on page 84](#)
- [Specifying Immediate-Leave Host Removal for MLD on page 85](#)
- [Filtering Unwanted MLD Reports at the MLD Interface Level on page 86](#)
- [Example: Modifying the MLD Robustness Variable on page 87](#)
- [Limiting the Maximum MLD Message Rate on page 88](#)
- [Enabling MLD Static Group Membership on page 89](#)
- [Example: Recording MLD Join and Leave Events on page 96](#)
- [Configuring the Number of MLD Multicast Group Joins on Logical Interfaces on page 98](#)
- [Disabling MLD on page 100](#)

Understanding MLD

The Multicast Listener Discovery (MLD) Protocol manages the membership of hosts and routers in multicast groups. IP version 6 (IPv6) multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol.


MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

In include mode, the receiver specifies the source or sources it is interested in receiving the multicast group traffic from. Exclude mode works the opposite of include mode. It allows the receiver to specify the source or sources it is not interested in receiving the multicast group traffic from.

For each attached network, a multicast routing device can be either a querier or a nonquerier. A querier routing device, usually one per subnet, solicits group membership information by transmitting MLD queries. When a host reports to the querier routing device that it has interested listeners, the querier routing device forwards the membership information to the rendezvous point (RP) routing device by means of the receiver's (host's) designated router (DR). This builds the rendezvous-point tree (RPT) connecting the host with interested listeners to the RP routing device. The RPT is the initial path used by the sender to transmit information to the interested listeners. Nonquerier routing devices do not transmit MLD queries on a subnet but can do so if the querier routing device fails.

All MLD-configured routing devices start as querier routing devices on each attached subnet (see [Figure 3 on page 76](#)). The querier routing device on the right is the receiver's DR.

Figure 9: Routing Devices Start Up on a Subnet




To elect the querier routing device, the routing devices exchange query messages containing their IPv6 source addresses. If a routing device hears a query message whose IPv6 source address is numerically lower than its own selected address, it becomes a nonquerier. In [Figure 4 on page 76](#), the routing device on the left has a source address numerically lower than the one on the right and therefore becomes the querier routing device.




NOTE: In the practical application of MLD, several routing devices on a subnet are nonqueriers. If the elected querier routing device fails, query messages are exchanged among the remaining routing devices. The routing device with the lowest IPv6 source address becomes the new querier routing device. The IPv6 Neighbor Discovery Protocol (NDP) implementation drops incoming Neighbor Announcement (NA) messages that have a broadcast or multicast address in the target link-layer address option. This behavior is recommended by RFC 2461.

Figure 10: Querier Routing Device Is Determined




The querier routing device sends general MLD queries on the **link-scope all-nodes** multicast address FF02::1 at short intervals to all attached subnets to solicit group membership information (see [Figure 5 on page 77](#)). Within the query message is the *maximum response delay* value, specifying the maximum allowed delay for the host to respond with a report message.

Figure 11: General Query Message Is Issued




If interested listeners are attached to the host receiving the query, the host sends a report containing the host's IPv6 address to the routing device (see [Figure 6 on page 77](#)). If the reported address is not yet in the routing device's list of multicast addresses with interested listeners, the address is added to the list and a timer is set for the address. If the address is already on the list, the timer is reset. The host's address is transmitted to the RP in the PIM domain.

Figure 12: Reports Are Received by the Querier Routing Device




If the host has no interested multicast listeners, it sends a done message to the querier routing device. On receipt, the querier routing device issues a multicast address-specific query containing the last **listener query interval** value to the multicast address of the host. If the routing device does not receive a report from the multicast address, it removes the multicast address from the list and notifies the RP in the PIM domain of its removal (see [Figure 7 on page 77](#)).

Figure 13: Host Has No Interested Receivers and Sends a Done Message to Routing Device



If a done message is not received by the querier routing device, the querier routing device continues to send multicast address-specific queries. If the timer set for the address on receipt of the last report expires, the querier routing device assumes there are no longer interested listeners on that subnet, removes the multicast address from the list, and notifies the RP in the PIM domain of its removal (see [Figure 8 on page 78](#)).

Figure 14: Host Address Timer Expires and Address Is Removed from Multicast Address List



Configuring MLD

To configure the Multicast Listener Discovery (MLD) Protocol, include the **ml**d statement:

```
ml {
  accounting;
  interface interface-name {
    disable;
    (accounting | no-accounting);
    group-policy [ policy-names ];
    immediate-leave;
    oif-map [ map-names ];
    passive;
    ssm-map ssm-map-name;
    static (Protocols MLD) {
      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;  
}
```

You can include this statement at the following hierarchy levels:

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

By default, MLD is enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or the Distance Vector Multicast Routing Protocol (DVMRP).

Enabling MLD

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

MLD specifies different behaviors for multicast listeners and for routers. When a router is also a listener, the router responds to its own messages. If a router has more than one interface to the same link, it needs to perform the router behavior over only one of those interfaces. Listeners, on the other hand, must perform the listener behavior on all interfaces connected to potential receivers of multicast traffic.

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

To explicitly enable MLD:

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

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

2. Check to see if MLD is disabled on any interfaces. In the following example, MLD is disabled on a Gigabit Ethernet interface.

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

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

```
[edit protocols mld]
```

```
delete interface ge-0/0/0.0 disable
```

4. Verify the configuration.

```
[edit protocols mld]
```

```
user@host# show
```

```
interface fe-0/0/0.0;
```

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

5. Verify the operation of MLD by checking the output of the **show mld interface** command.

Modifying the MLD Version

By default, the router supports MLD version 1 (MLDv1). To enable the router to use MLD version 2 (MLDv2) for source-specific multicast (SSM) only, include the **version 2** statement.

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

If a source address is specified in a multicast group that is statically configured, the version must be set to MLDv2.

To change an MLD interface to version 2:

1. Configure the MLD interface.

```
[edit protocols mld]
```

```
user@host# set interface fe-0/0/0.0 version 2
```

2. Verify the configuration by checking the **version** field in the output of the **show mld interface** command. The **show mld statistics** command has version-specific output fields, such as the counters in the **MLD Message type** field.

Modifying the MLD Host-Query Message Interval

The objective of MLD is to keep routers up to date with IPv6 group membership of the entire subnet. Routers need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one router is elected the querier. The MLD querier router periodically sends general host-query messages on each attached network to solicit membership information. These messages solicit group membership information and are sent to the **link-scope all-nodes** address **FF02::1**. A general host-query message has a maximum response time that you can set by configuring the query response interval.

The query response timeout, the query interval, and the robustness variable are related in that they are all variables that are used to calculate the multicast listener interval. The multicast listener interval is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The multicast listener interval is calculated as the (robustness variable x query-interval) + (1 x query-response-interval). If no reports are received for a particular group before the multicast listener interval has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

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

To modify the query interval:

1. Configure the interval.

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

The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the **MLD Query Interval** field in the output of the **show mld interface** command.
3. Verify the operation of the query interval by checking the **Listener Query** field in the output of the **show mld statistics** command.

Modifying the MLD Query Response Interval

The query response interval is the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. 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.

The query response timeout, the query interval, and the robustness variable are related in that they are all variables that are used to calculate the multicast listener interval. The multicast listener interval is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The multicast listener interval is calculated as the (robustness variable x query-interval) + (1 x query-response-interval). If no reports are received for a particular group before the multicast listener interval has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

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

To modify the query response interval:

1. Configure the interval.

```
[edit protocols mld]  
user@host# set query-response-interval 0.5
```

2. Verify the configuration by checking the **MLD Query Response Interval** field in the output of the **show mld interface** command.
3. Verify the operation of the query interval by checking the **Listener Query** field in the output of the **show mld statistics** command.

Modifying the MLD Last-Member Query Interval

The last-member query interval (also called the last-listener query interval) is the maximum amount of time between group-specific query messages, including those sent

in response to done messages sent on the **link-scope-all-routers** address FF02::2. You can lower this interval to reduce the amount of time it takes a router to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group (done) message from a host, the routing device sends multiple group-specific queries to the group. The querier sends a specific number of these queries, and it sends them at a specific interval. The number of queries sent is called the last-listener query count. The interval at which the queries are sent is called the last-listener query interval. Both settings are configurable, thus allowing you to adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

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

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

To modify this interval:

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

```
[edit protocols mld]
user@host# set query-last-member-interval 0.1
```
2. Verify the configuration by checking the **MLD Last Member Query Interval** field in the output of the **show igmp interfaces** command.



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

Specifying Immediate-Leave Host Removal for MLD

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

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

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

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both MLD version 1 and MLD version 2.



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

To enable immediate leave:

1. Configure immediate leave on the MLD interface.

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

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

Filtering Unwanted MLD Reports at the MLD Interface Level

Suppose you need to limit the subnets that can join a certain multicast group. The **group-policy** statement enables you to filter unwanted MLD reports at the interface level.

When the **group-policy** statement is enabled on a router, after the router receives an MLD report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only MLD group addresses (for MLDv1) by using the policy's **route-filter** statement to match the group address. You define the policy to match MLD (source, group) addresses (for MLDv2) by using the policy's **route-filter** statement to match the group address and the policy's **source-address-filter** statement to match the source address.

To filter unwanted MLD reports:

1. Configure an MLDv1 policy.

```
[edit policy-statement reject_policy_v1]
user@host# set from route-filter fec0:1:1:4::/64 exact
user@host# set then reject
```

2. Configure an MLDv2 policy.

```
[edit policy-statement reject_policy_v2]
```

```

user@host# set from route-filter fec0:1:1:4::/64 exact
user@host# set from source-address-filter fe80::2e0:81ff:fe05:1a8d/32 orlonger
user@host# set then reject

```

3. Apply the policies to the MLD interfaces where you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running MLDv1 and **ge-0/1/1.0** is running MLDv2.

```

[edit protocols mld]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v1
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v2

```

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

Example: Modifying the MLD Robustness Variable

This example shows how to configure and verify the MLD robustness variable in a multicast domain.

- [Requirements on page 87](#)
- [Overview on page 87](#)
- [Configuration on page 88](#)
- [Verification on page 88](#)

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable IPv6 unicast routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable PIM. See “[PIM Overview](#)” on page 125.

Overview

The MLD robustness variable can be fine-tuned to allow for expected packet loss on a subnet. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

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

- Group member interval—Amount of time that must pass before a multicast router determines that there are no more members of a group on a network. This interval is calculated as follows: (robustness variable x query-interval) + (1 x query-response-interval).

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

By default, the robustness variable is set to 2. The number can be from 2 through 10. You might want to increase this value if you expect a subnet to lose packets.

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, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols mld robust-count 5
```

Step-by-Step Procedure

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

To change the value of the robustness variable:

1. Configure the robust count.

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

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

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

Verification

To verify the configuration is working properly, check the **MLD Robustness Count** field in the output of the **show mld interfaces** command.

Limiting the Maximum MLD Message Rate

You can change the limit for the maximum number of MLD packets transmitted in 1 second by the router.

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

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

Enabling MLD Static Group Membership

- [Create a MLD Static Group Member on page 89](#)
- [Automatically create static groups on page 90](#)
- [Automatically increment group addresses on page 91](#)
- [Specify multicast source address \(in SSM mode\) on page 92](#)
- [Automatically specify multicast sources on page 93](#)
- [Automatically increment source addresses on page 94](#)
- [Exclude multicast source addresses \(in SSM mode\) on page 95](#)

Create a MLD Static Group Member

You can create MLD static group membership to test multicast forwarding without a receiver host. When you enable MLD static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts.

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

When you configure static groups on an interface on which you want to receive multicast traffic, you can specify the number of static groups to be automatically created.

In this example, you create static group ff0e::1:ff05:1a8d.

1. Configure the static groups to be created by including the **static** statement and **group** statement and specifying which IPv6 multicast address of the group to be created.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d;
  }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show mld group** command to verify that static group ff0e::1:ff05:1a8d has been created.

```
user@host> show mld group
Interface: fe-0/1/2
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
```



NOTE: You must specify a unique address for each group.

Automatically create static groups

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

In this example, you create three groups.

1. Configure the number of static groups to be created by including the **group-count** statement and specifying the number of groups to be created.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
group-count 3
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static groups ff0e::1:ff05:1a8d, ff0e::1:ff05:1a8e, and ff0e::1:ff05:1a8f have been created.

```
user@host> show mld group

Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8e
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8f
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
```

Automatically increment group addresses

When you configure static groups on an interface on which you want to receive multicast traffic and you specify the number of static groups to be automatically created, you can also configure the group address to be automatically incremented by some number of addresses.

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

1. Configure the group address increment by including the **group-increment** statement and specifying the number by which the address should be incremented for each group. The increment is specified in a format similar to an IPv6 address.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
group-count 3 group-increment ::2
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      group-increment ::2;
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static groups ff0e::1:ff05:1a8d, ff0e::1:ff05:1a8f, and ff0e::1:ff05:1a91 have been created.

```
user@host> show mld group

Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8f
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a91
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
```

Specify multicast source address (in SSM mode)

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify the multicast source address to be accepted.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the MLD version must be set to MLDv2 on the interface. MLDv1 is the default value.

In this example, you create group ff0e::1:ff05:1a8d and accept IPv6 address fe80::2e0:81ff:fe05:1a8d as the only source.

1. Configure the source address by including the **source** statement and specifying the IPv6 address of the source host.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
source fe80::2e0:81ff:fe05:1a8d
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      source fe80::2e0:81ff:fe05:1a8d;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group ff0e::1:ff05:1a8d has been created and that source fe80::2e0:81ff:fe05:1a8d has been accepted.

```
user@host> show mld group

Interface: fe-0/1/2
Group: ff0e::1:ff05:1a8d
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
```


Automatically specify multicast sources

When you configure static groups on an interface on which you want to receive multicast traffic, you can specify a number of multicast sources to be automatically accepted.

In this example, you create static group ff0e::1:ff05:1a8d and accept fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8e, and fe80::2e0:81ff:fe05:1a8f as the source addresses.

1. Configure the number of multicast source addresses to be accepted by including the **source-count** statement and specifying the number of sources to be accepted.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
source fe80::2e0:81ff:fe05:1a8d source-count 3
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      source fe80::2e0:81ff:fe05:1a8d {
        source-count 3;
      }
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group ff0e::1:ff05:1a8d has been created and that sources fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8e, and fe80::2e0:81ff:fe05:1a8f have been accepted.

```
user@host> show mld group

Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8e
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8f
    Last reported by: Local
    Timeout: 0 Type: Static
```

Automatically increment source addresses

When you configure static groups on an interface on which you want to receive multicast traffic, and specify a number of multicast sources to be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted.

In this example, you create static group ff0e::1:ff05:1a8d and accept fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8f, and fe80::2e0:81ff:fe05:1a91 as the sources.

1. Configure the number of multicast source addresses to be accepted by including the **source-increment** statement and specifying the number of sources to be accepted.

```
[edit protocols mld]
```

```
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
source fe80::2e0:81ff:fe05:1a8d source-count 3 source-increment ::2
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld
```

```
interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      source fe80::2e0:81ff:fe05:1a8d {
        source-count 3;
        source-increment ::2;
      }
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group ff0e::1:ff05:1a8d has been created and that sources fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8f, and fe80::2e0:81ff:fe05:1a91 have been accepted.

```
user@host> show mld group
```

```
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8f
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e2::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a91
    Last reported by: Local
    Timeout: 0 Type: Static
```

```

Interface: fe-0/1/2
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8f
Last reported by: Local
Timeout: 0 Type: Static
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a91
Last reported by: Local
Timeout: 0 Type: Static

```

Exclude multicast source addresses (in SSM mode)

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the configured source address. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the configured source address.

If a source address is specified in a multicast group that is statically configured, the MLD version must be set to MLDv2 on the interface. MLDv1 is the default value.

In this example, you exclude address fe80::2e0:81ff:fe05:1a8d as a source for group ff0e::1:ff05:1a8d.

1. Configure a multicast static group to operate in exclude mode by including the **exclude** statement and specifying which IPv6 source address to be excluded.

```

[edit protocols mld]
user@host# set interface fe-0/1/2 static (Protocols MLD) group ff0e::1:ff05:1a8d
exclude source fe80::2e0:81ff:fe05:1a8d

```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```

user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      exclude;
      source fe80::2e0:81ff:fe05:1a8d;
    }
  }
}

```

- After you have committed the configuration and the source is sending traffic, use the **show mld group detail** command to verify that static group ff0e::1:ff05:1a8d has been created and that the static group is operating in exclude mode.

```
user@host> show mld group detail
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Group mode: Exclude
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
```

Similar configuration is available for IPv4 multicast traffic using the IGMP protocol.

Example: Recording MLD Join and Leave Events

This example shows how to determine whether MLD tuning is needed in a network by configuring the routing device to record MLD join and leave events.

- [Requirements on page 96](#)
- [Overview on page 96](#)
- [Configuration on page 97](#)
- [Verification on page 98](#)

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable IPv6 unicast routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable PIM. See “PIM Overview” on page 125.

Overview

[Table 7 on page 96](#) describes the recordable MLD join and leave events.

Table 7: MLD Event Messages

ERRMSG Tag	Definition
RPD_MLD_JOIN	Records MLD join events.
RPD_MLD_LEAVE	Records MLD leave events.
RPD_MLD_ACCOUNTING_ON	Records when MLD accounting is enabled on an MLD interface.
RPD_MLD_ACCOUNTING_OFF	Records when MLD accounting is disabled on an MLD interface.

Table 7: MLD Event Messages (*continued*)

ERRMSG Tag	Definition
RPD_MLD_MEMBERSHIP_TIMEOUT	Records MLD membership timeout events.

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, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols mld interface fe-0/1/0.2 accounting
set system syslog file mld-events any info
set system syslog file mld-events match ".*RPD_MLD_JOIN.* | .*RPD_MLD_LEAVE.* |
.*RPD_MLD_ACCOUNTING.* | .*RPD_MLD_MEMBERSHIP_TIMEOUT.*"
set system syslog file mld-events archive size 100000
set system syslog file mld-events archive files 3
set system syslog file mld-events archive transfer-interval 1440
set system syslog file mld-events archive archive-sites "ftp://user@host1//var/tmp"
password "anonymous"
set system syslog file mld-events archive archive-sites "ftp://user@host2//var/tmp"
password "test"
```

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

To configure recording of MLD join and leave events:

1. Enable accounting globally or on an MLD interface. This example shows the interface configuration.

```
[edit protocols mld]
user@host# set interface fe-0/1/0.2 accounting
```

2. Configure the events to be recorded, and filter the events to a system log file with a descriptive filename, such as **mld-events**.

```
[edit system syslog file mld-events]
user@host# set any info
[edit system syslog file mld-events]
user@host# set match ".*RPD_MLD_JOIN.* | .*RPD_MLD_LEAVE.* |
.*RPD_MLD_ACCOUNTING.* | .*RPD_MLD_MEMBERSHIP_TIMEOUT.*"
```

3. Periodically archive the log file.

This example rotates the file every 24 hours (1440 minutes) when it reaches 100 KB and keeps three files.

```
[edit system syslog file mld-events]
user@host# set archive size 100000
[edit system syslog file mld-events]
user@host# set archive files 3
[edit system syslog file mld-events]
```

```
user@host# set archive archive-sites "ftp://user@host1//var/tmp" password
"anonymous"
[edit system syslog file mld-events]
user@host# set archive archive-sites "ftp://user@host2//var/tmp" password "test"
[edit system syslog file mld-events]
user@host# set archive transfer-interval 1440
[edit system syslog file mld-events]
user@host# set archive start-time 2011-01-07:12:30
```

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

```
[edit system syslog file mld-events]]
user@host# commit
```

Verification

You can view the system log file by running the **file show** command.

```
user@host> file show mld-events
```

You can monitor the system log file as entries are added to the file by running the **monitor start** and **monitor stop** commands.

```
user@host> monitor start mld-events
```

```
*** mld-events ***
Apr 16 13:08:23 host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command 'run
monitor start mld-events '
monitor
```

Configuring the Number of MLD Multicast Group Joins on Logical Interfaces

The **group-limit** statement enables you to limit the number of MLD multicast group joins for logical interfaces. When this statement is enabled on a router running MLD version 2, the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for MLD multicast groups, keep the following in mind:

- Each any-source group (*G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in MLDv2 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The

groups must then request to rejoin the network (up to the newly configured group limit).

- You can dynamically limit multicast groups on MLD logical interfaces by using dynamic profiles. For detailed information about creating dynamic profiles, see the *Junos OS Broadband Subscriber Management and Services Library*.

Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for MLD multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of MLD multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The **group-threshold** statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a percentage of the group limit, so you must configure the **group-limit** statement to configure a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs a warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the **log-interval** statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for MLD multicast group joins.

To limit multicast group joins on an MLD logical interface:

1. Access the logical interface at the MLD protocol hierarchy level.

```
[edit]
user@host# edit protocols mld interface interface-name
```

2. Specify the group limit for the interface.

```
[edit protocols mld interface interface-name]
user@host# set group-limit limit
```

3. (Optional) Configure the threshold at which a warning message is logged.

```
[edit protocols mld interface interface-name]
user@host# set group-threshold value
```

4. (Optional) Configure the amount of time between log messages.

```
[edit protocols mld interface interface-name]
user@host# set log-interval seconds
```

To confirm your configuration, use the **show protocols mld** command. To verify the operation of MLD on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the **show mld interface** command.

Disabling MLD

To disable MLD on an interface, include the **disable** statement:

```
interface interface-name {  
    disable;  
}
```

You can include this statement at the following hierarchy levels:

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

Related Documentation

- *Configuring IGMP*

CHAPTER 5

Using MLD Snooping

- [Understanding MLD Snooping on page 101](#)
- [Configuring MLD Snooping on a VLAN \(CLI Procedure\) on page 109](#)
- [Example: Configuring MLD Snooping on page 115](#)
- [Verifying MLD Snooping on page 118](#)

Understanding MLD Snooping



NOTE: This overview uses Junos OS for switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Understanding MLD Snooping*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

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 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 102](#)
- [MLD Message Types on page 103](#)
- [How Hosts Join and Leave Multicast Groups on page 103](#)
- [Support for MLDv2 Multicast Sources on page 104](#)
- [MLD Snooping and Forwarding Interfaces on page 104](#)

- [General Forwarding Rules on page 105](#)
- [Examples of MLD Snooping Multicast Forwarding on page 105](#)


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 15 on page 102](#) shows an example of multicast traffic flow with MLD snooping enabled.

Figure 15: 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.

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, 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 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 105](#)
- [Scenario 2: Switch Forwarding Multicast Traffic to Another Switch on page 106](#)
- [Scenario 3: Switch Connected to Hosts Only \(No MLD Querier\) on page 107](#)
- [Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs on page 108](#)

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

In the topology shown in [Figure 16 on page 106](#), 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 16: Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts



g041195

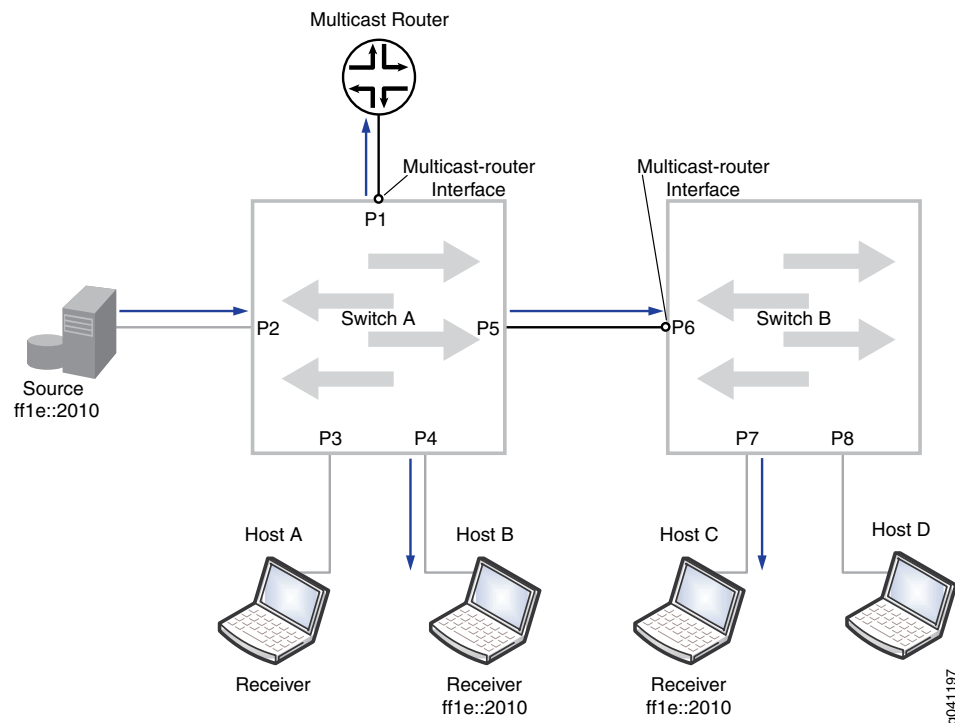
Scenario 2: Switch Forwarding Multicast Traffic to Another Switch

In the topology shown in [Figure 17 on page 107](#), 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 17: 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 18 on page 108](#), 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 18: Scenario 3: Switch Connected to Hosts Only (No MLD Querier)

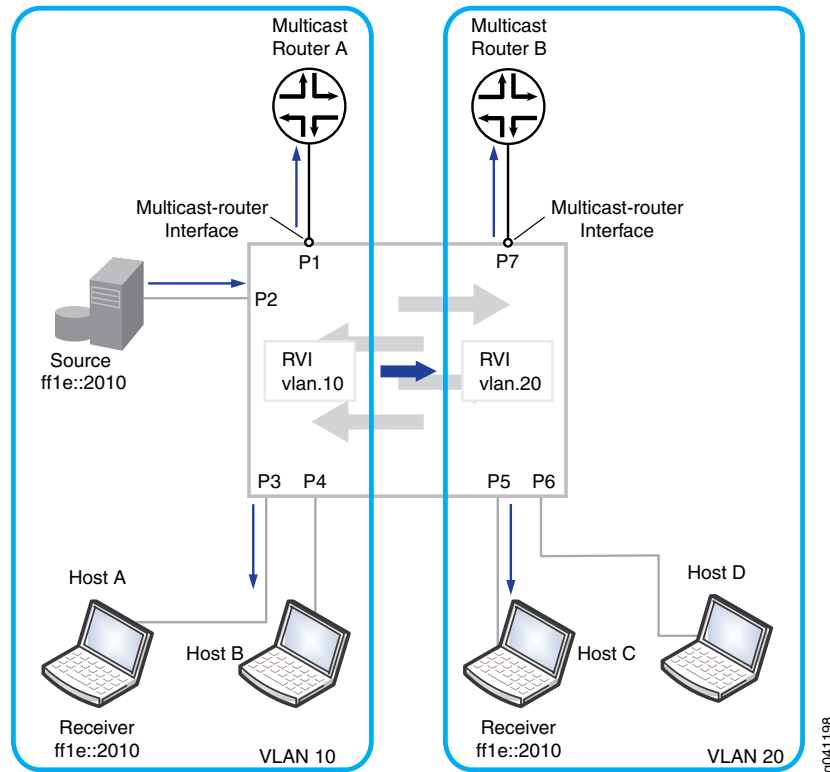


Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs

In the topology shown in [Figure 19 on page 109](#), 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.

Figure 19: Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs



- Related Documentation**
- *Example: Configuring MLD Snooping*
 - *Configuring MLD Snooping on a VLAN (CLI Procedure)*
 - *Verifying MLD Snooping (CLI Procedure)*

Configuring MLD Snooping on a VLAN (CLI Procedure)



NOTE: This task uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring MLD Snooping on a VLAN (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

You can enable MLD snooping on a VLAN to constrain the flooding of IPv6 multicast traffic on the 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.

You can perform the following configurations for each VLAN:

- Selectively enable MLD snooping on specific VLANs.
- Specify the MLD version for the general query that the switch sends on an interface when the interface comes up.
- Enable immediate leave to reduce 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 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.

This topic covers:

- [Enabling or Disabling MLD Snooping on VLANs on page 110](#)
- [Configuring the MLD Version on page 111](#)
- [Enabling Immediate Leave on page 111](#)
- [Configuring an Interface as a Multicast-Router Interface on page 112](#)
- [Configuring Static Group Membership on an Interface on page 113](#)
- [Changing the Timer and Counter Values on page 114](#)

Enabling or Disabling MLD Snooping on VLANs

MLD snooping is not enabled on any VLAN by default. You must explicitly enable MLD snooping on specific interfaces.

- To enable MLD snooping on a specific VLAN:

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



NOTE: You cannot enable 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 disable MLD snooping on a specific VLAN:

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

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 an interface:

```
[edit protocols]
user@switch# set mld interface interface-name version number
```

For example, to set the MLD version to version 2 on interface ge-0/0/2:

```
[edit protocols]
user@switch# set mld interface ge-0/0/2 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
```

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.

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 VLAN employee:

```
[edit protocols]
user@switch# set mld-snooping vlan employee 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 employee 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 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 IPv6 multicast deployments.

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 in seconds 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.

To configure the MLD query interval:

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

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

To configure the MLD query response interval:

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

To configure the MLD query last member interval:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan-name 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 anticipated packet loss.

For MLD snooping, you can configure **robust-count** for a specific VLAN. If a VLAN does not have **robust-count** configured, the 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** value by the **robust-count** value and then adding the **query-response-interval** to the product:

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

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

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

To display the time remaining in the multicast listener interval before a group times out, use the **show mld-snooping membership** command.

Related Documentation

- [Example: Configuring MLD Snooping on page 115](#)
- [Examples: Configuring MLD on page 78](#)
- [Verifying MLD Snooping on page 118](#)

Example: Configuring MLD Snooping



NOTE: This example uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

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. On the basis of 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 116](#)
- [Overview and Topology on page 116](#)
- [Configuration on page 117](#)
- [Verifying MLD Snooping Configuration on page 118](#)

Requirements

This example uses the following software and hardware components:

- One switch running Junos OS with ELS
- Junos OS Release 13.3 or later for EX Series switches or Junos OS Release 15.1X53-D10 or later for QFX10000 switches

Before you configure MLD snooping, be sure you have:

- Configured the vlan 100 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)* or *Configuring VLANs*.

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 topology for this example is illustrated in [Figure 20 on page 116](#).

Figure 20: MLD Snooping Topology Example



In this sample 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 the 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 118](#)

Verifying MLD Snooping Interface Membership on VLAN vlan100

Purpose Verify that MLD snooping is enabled on the VLAN vlan 100 and that the multicast-router interface is statically configured:

Action Show the MLD snooping information for ge-0/0/12.0:

```
user@switch> show mld snooping interface
Instance: default-switch
```

```
Vlan: vlan100
```

```
Learning-Domain: default
Interface: ge-0/0/12.0
  State:          Up Groups:      3
  Immediate leave: On
  Router interface: yes
```

```
Configured Parameters:
MLD Query Interval: 125.0
MLD Query Response Interval: 10.0
MLD Last Member Query Interval: 1.0
MLD Robustness Count: 2
```

Meaning MLD snooping is running on **vlan100**, and interface **ge-0/0/12.0** is a statically configured multicast-router interface. Immediate leave is enabled on the interface.

Related Documentation

- [Configuring MLD Snooping on a VLAN \(CLI Procedure\) on page 109](#)
- [Verifying MLD Snooping on page 118](#)
- [Understanding MLD Snooping on page 101](#)

Verifying MLD Snooping



NOTE: This topic uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Verifying MLD Snooping (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

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

It covers:

- [Verifying MLD Snooping Memberships on page 119](#)
- [Verifying MLD Snooping Interfaces on page 119](#)
- [Viewing MLD Snooping Statistics on page 120](#)
- [Viewing MLD Snooping Routing Information on page 121](#)

Verifying MLD Snooping Memberships

Purpose Verify that MLD snooping is enabled on a VLAN and determine group memberships.

Action Enter the following command:

```
user@switch> show mld snooping membership detail
Instance: default-switch
```

```
Vlan: v1
```

```
Learning-Domain: default
Interface: ge-0/0/1.0, Groups: 1
  Group: ff05::1
    Group mode: Exclude
    Source: ::
    Last reported by: fe80::
    Group timeout: 259 Type: Dynamic
Interface: ge-0/0/2.0, Groups: 0
```

Meaning The switch has multicast membership information for one VLAN on the switch, **v1**. 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 about the group memberships for the VLAN:
 - Currently, the VLAN has membership in only one multicast group, **ff05::1**.
 - The host or hosts that have reported membership in the group are on interface **ge-0/0/1.0**.
 - The last host that reported membership in the group has address **fe80::**.
 - The interface group membership will time out in **259** seconds if no hosts respond to membership queries during this interval.
 - The group membership has been learned by MLD snooping, as indicated by **Dynamic**.

Verifying MLD Snooping Interfaces

Purpose Display MLD snooping information for each interface on which MLD snooping is enabled.

Action Enter the following command:

```
user@switch> show mld snooping interface
Instance: default-switch
```

```
Vlan: v100
```

```

Learning-Domain: default
Interface: ge-0/0/1.0
  State:          Up Groups:      1
  Immediate leave: Off
  Router interface: no
Interface: ge-0/0/2.0
  State:          Up Groups:      0
  Immediate leave: Off
  Router interface: no

Configured Parameters:
MLD Query Interval: 125.0
MLD Query Response Interval: 10.0
MLD Last Member Query Interval: 1.0
MLD Robustness Count: 2

```

Meaning MLD snooping is configured on one VLAN on the switch, **v100**. Each interface in each VLAN is listed and the following information is provided:

- How many multicast groups the interface belongs to.
- Whether immediate leave has been configured for the interface.
- Whether the interface is a multicast-router interface.

The output also shows the configured parameters for the MLD querier.

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
Vlan: v1
MLD Message type      Received      Sent  Rx errors
Listener Query (v1/v2)      0            4      0
Listener Report (v1)      447          0      0
Listener Done (v1/v2)      0            0      0
Listener Report (v2)       0            0      0
Other Unknown types              0
Vlan: v2
MLD Message type      Received      Sent  Rx errors
Listener Query (v1/v2)      0            4      0
Listener Report (v1)      154          0      0
Listener Done (v1/v2)      0            0      0
Listener Report (v2)       0            0      0
Other Unknown types              0
Instance: default-switch
MLD Message type      Received      Sent  Rx errors
Listener Query (v1/v2)      0            8      0
Listener Report (v1)      601          0      0
Listener Done (v1/v2)      0            0      0
Listener Report (v2)       0            0      0
Other Unknown types              0

```

```

MLD Global Statistics
Bad Length           0
Bad Checksum         0
Bad Receive If       0
Rx non-local         0
Timed out            0

```

Meaning The output shows how many MLD messages of each type—**Queries**, **Done**, **Report**—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 **Rx 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 **Other Unknown types**.

Viewing MLD Snooping Routing Information

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

Action Enter the following command:

```

user@switch>show multicast snooping route
Nexthop Bulking: OFF

```

```

Family: INET6

```

```

Group: ff00::/8
Source: ::/128
Vlan: v1

Group: ff02::1/128
Source: ::/128
Vlan: v1
Downstream interface list:
ge-1/0/16.0

```

```

Group: ff05::1/128
Source: ::/128
Vlan: v1
Downstream interface list:
ge-1/0/16.0

```

```

Group: ff06::1/128
Source: ::/128
Vlan: v1
Downstream interface list:
ge-1/0/16.0

```

Meaning The output shows the next-hop interfaces for a given multicast group on a VLAN. For example, route **ff02::1/128** on VLAN **v1** has the next-hop interface **ge-1/0/16.0**.

Related Documentation

- *clear mld snooping membership*
- *clear mld snooping statistics*
- [Example: Configuring MLD Snooping on page 115](#)

- [Configuring MLD Snooping on a VLAN \(CLI Procedure\) on page 109](#)

PART 2

Configuring PIM

- [Using PIM Basic Features on page 125](#)
- [Using PIM Sparse Mode on page 143](#)
- [Using PIM Dense Mode and PIM Sparse-Dense Mode on page 157](#)
- [Using Source-Specific Multicast on page 163](#)
- [Using Static RP on page 181](#)
- [Using Anycast RP on page 185](#)
- [Using Auto-RP on page 195](#)
- [Using PIM Bootstrap Router on page 201](#)
- [Using PIM Filtering on page 205](#)
- [Using PIM RPT and SPT Cutover on page 213](#)

CHAPTER 6

Using PIM Basic Features

- [PIM Overview on page 125](#)
- [PIM on Aggregated Interfaces on page 128](#)
- [Changing the PIM Version on page 128](#)
- [Modifying the PIM Hello Interval on page 128](#)
- [Preserving Multicast Performance by Disabling Response to the ping Utility on page 129](#)
- [Configuring PIM Trace Options on page 130](#)
- [Configuring Interface Priority for PIM Designated Router Selection on page 132](#)
- [Configuring PIM Designated Router Election on Point-to-Point Links on page 133](#)
- [Configuring BFD for PIM on page 134](#)
- [Configuring BFD Authentication for PIM on page 135](#)
- [Disabling PIM on page 139](#)

PIM Overview

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

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

PIM is one of the most rapidly evolving specifications on the Internet today. Since its introduction in 1995, PIM has already seen two major revisions to its packet structure (PIM version 1 [PIMv1] and PIM version 2 [PIMv2]), two major RFCs (RFC 2362 obsoleted RFC 2117), and numerous drafts describing major components of PIM, such as many-to-many trees and source-specific multicast (SSM). Long-lasting RFCs are not a feature of PIM, and virtually all of PIM must be researched, understood, and implemented

directly from Internet drafts. In fact, no current RFC describes PIMv1 at all. The drafts have all expired, and PIMv1 was never issued as an official RFC.

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

PIMv1 and PIMv2 can coexist on the same router and even on the same interface. The main difference between PIMv1 and PIMv2 is the packet format. PIMv1 messages use Internet Group Management Protocol (IGMP) packets, whereas PIMv2 has its own IP protocol number (103) and packet structure. All routers connecting to an IP subnet such as a LAN must use the same PIM version. Some PIM implementations can recognize PIMv1 packets and automatically switch the router interface to PIMv1. Because the difference between PIMv1 and PIMv2 involves the message format, but not the meaning of the message or how the router processes the PIM message, a router can easily mix PIMv1 and PIMv2 interfaces.

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

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

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

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

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

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

- Bidirectional PIM is similar to sparse mode, and is especially suited to applications that must scale to support a large number of dispersed sources and receivers. In bidirectional PIM, routers build shared bidirectional trees and do not switch to a source-based tree.

Bidirectional PIM scales well because it needs no source-specific (S,G) state. Instead, it builds only group-specific (*G) state.

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

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

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



NOTE: On SRX Series devices, PIM does not support upstream and downstream interfaces across different virtual routers in flow mode.

Basic PIM Network Components

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

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

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

The bootstrap router does not have to be an RP itself, although this is a common implementation. The bootstrap router's main function is to manage the collection of RPs and allow interested receivers to find the source of their group's multicast traffic.

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

Bidirectional PIM RPs, unlike RPs for PIM sparse mode, do not need to perform PIM Register tunneling or other specific protocol action. Bidirectional PIM RPs implement no specific functionality. RP addresses are simply a location in the network to rendezvous toward. In fact, for bidirectional PIM, RP addresses need not be loopback interface addresses or even be addresses configured on any router, as long as they are covered by a subnet that is connected to a bidirectional PIM-capable router and advertised to the network.

Related Documentation

- *Supported IP Multicast Protocol Standards*

PIM on Aggregated Interfaces

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

Related Documentation

- [PIM Overview on page 125](#)
- [interface on page 353](#)

Changing the PIM Version

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

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

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

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

Modifying the PIM Hello Interval

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

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

the neighbor. In the case of a 30-second interval, the amount of time a routing device waits for a response is 105 seconds.

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

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

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

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

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

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

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

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

Related Documentation

- [show pim neighbors on page 541](#)

Preserving Multicast Performance by Disabling Response to the ping Utility

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

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

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

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

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

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

```
user@host> show system statistics icmp

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

- Related Documentation**
- *Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping Packets*
 - *show system statistics icmp*

Configuring PIM Trace Options

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

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

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

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

To configure tracing operations for PIM:

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

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

2. Configure the filename for the PIM trace file.

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

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

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

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

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

5. (Optional) Enable unrestricted file access.

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

6. Configure tracing flags.

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

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

7. View the trace file.

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

- Related Documentation**
- [PIM Overview on page 125](#)
 - *Tracing and Logging Junos OS Operations*

Configuring Interface Priority for PIM Designated Router Selection

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

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

To configure the interface designated router priority:

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

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

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


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

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

```
Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

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

- Related Documentation**
- [Configuring PIM Designated Router Election on Point-to-Point Links on page 133](#)
 - [Understanding PIM Sparse Mode on page 143](#)
 - [show pim neighbors on page 541](#)

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

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

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

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

```
[edit routing-instances PIM.master protocols pim]
user@host# set dr-election-on-p2p
```
2. Verify the configuration by checking the **State** field in the output of the **show pim interfaces** command. The possible values for the **State** field are DR, NotDR, and P2P. When a point-to-point link interface is elected to be the DR, the interface state becomes DR instead of P2P.
3. If the **show pim interfaces** command continues to report the P2P state, consider running the **restart routing** command on both routers on the point-to-point link. Then recheck the state.



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

[edit]

user@host# run restart routing

**Related
Documentation**

- [Understanding PIM Sparse Mode on page 143](#)
- [Configuring Interface Priority for PIM Designated Router Selection on page 132](#)
- [show pim interfaces on page 517](#)

Configuring BFD for PIM

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

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

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

To enable failure detection:

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

This example shows the global configuration.

[edit protocols pim]

user@host# edit interface fe-1/0/0.0 family inet **bfd-liveness-detection**

2. Configure the minimum transmit interval.

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

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

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

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

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

4. (Optional) Configure other BFD settings.

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

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

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

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

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

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

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

7. Configure the BFD version.

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

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

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

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

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

Related Documentation

- *show bfd session*

Configuring BFD Authentication for PIM

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

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

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

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

- [Configuring BFD Authentication Parameters on page 136](#)
- [Viewing Authentication Information for BFD Sessions on page 137](#)

Configuring BFD Authentication Parameters

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

To configure BFD authentication:

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

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



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

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

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

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



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

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

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

```
[edit security]
```

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

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

```
[edit protocols pim]
```

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

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

Viewing Authentication Information for BFD Sessions

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

The following example shows BFD authentication configured for the **ge-0/1/5** interface. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-pim**. The authentication keychain is configured with two keys. Key 1 contains the secret data “\$9\$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm” and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data “\$9\$a5jiKW9L.reP38ny.TszF2/9” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols pim]
```

```
interface ge-0/1/5 {
  family inet {
    bfd-liveness-detection {
      authentication {
        key-chain bfd-pim;
        algorithm keyed-sha-1;
      }
    }
  }
}
```

```
[edit security]
```

```
authentication key-chains {
  key-chain bfd-pim {
    key 1 {
      secret "$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm";
      start-time "2009-6-1.09:46:02 -0700";
    }
  }
}
```

```

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

```

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

show bfd session detail

```
user@host# show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
50.0.0.2	Up	ge-0/1/5.0	0.900	0.300	3

Client PIM, TX interval 0.300, RX interval 0.300, **Authenticate**
 Session up time 3d 00:34
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated

show bfd session extensive

```
user@host# show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
50.0.0.2	Up	ge-0/1/5.0	0.900	0.300	3

Client PIM, TX interval 0.300, RX interval 0.300, **Authenticate**
keychain bfd-pim, algo keyed-sha-1, mode strict
 Session up time 00:04:42
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated
 Min async interval 0.300, min slow interval 1.000
 Adaptive async TX interval 0.300, RX interval 0.300
 Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3
 Remote min TX interval 0.300, min RX interval 0.300, multiplier 3
 Local discriminator 2, remote discriminator 2
 Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-pim, algo keyed-sha-1, mode strict

Related Documentation

- [Understanding Bidirectional Forwarding Detection Authentication for PIM](#)
- [Configuring BFD for PIM on page 134](#)
- [authentication-key-chains](#)
- [bfd-liveness-detection on page 331](#)
- [show bfd session](#)

Disabling PIM

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

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

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

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

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

- [Disabling the PIM Protocol on page 139](#)
- [Disabling PIM on an Interface on page 140](#)
- [Disabling PIM for a Family on page 140](#)
- [Disabling PIM for a Rendezvous Point on page 141](#)

Disabling the PIM Protocol

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

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

To disable the PIM protocol:

1. Include the **disable** statement.

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

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

Disabling PIM on an Interface

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

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

To disable PIM on an interface:

1. Include the **disable** statement.

```
user@host# set protocols pim interface fe-0/1/0 disable
```

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

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

Disabling PIM for a Family

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

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

To disable PIM for a family:

1. Include the **disable** statement.

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

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

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

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


Disabling PIM for a Rendezvous Point

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

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

To disable PIM for an RP family:

1. Use the **disable** statement.

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

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

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


CHAPTER 7

Using PIM Sparse Mode

- [Understanding PIM Sparse Mode on page 143](#)
- [Designated Router on page 146](#)
- [Enabling PIM Sparse Mode on page 146](#)
- [Configuring PIM Join Load Balancing on page 147](#)
- [Modifying the Join State Timeout on page 151](#)
- [Example: Enabling Join Suppression on page 151](#)

Understanding PIM Sparse Mode

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



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

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

If the RP router has interested receivers in the PIM sparse-mode domain, it sends a PIM join message toward the source to build a shortest-path tree (SPT) back to the source. The source sends multicast packets out on the LAN, and the source DR encapsulates

the packets in a PIM register message and forwards the message toward the RP router by means of unicast. The RP router receives PIM register messages back from the source, and thus adds a new source to the distribution tree, keeping track of sources in a PIM table. Once an RP router receives packets natively (with S,G), it sends a register stop message to stop receiving the register messages by means of unicast.

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

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



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

The major characteristics of PIM sparse mode are as follows:

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

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

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

- Sources must be able to send to all RPs.
- RPs must all know one another.
- Receivers must send explicit join messages to a known RP.
- Receivers initially need to know only one RP (they later learn about others).
- Receivers can explicitly prune themselves from a tree.
- Receivers that never transition to a source-based tree are effectively running Core Based Trees (CBT).


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

Rendezvous Point

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

The RP router is downstream from the source and forms one end of the shortest-path tree. As shown in [Figure 21 on page 145](#), the RP router is upstream from the receiver and thus forms one end of the rendezvous-point tree.

Figure 21: Rendezvous Point As Part of the RPT and SPT



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

RP Mapping Options

RPs can be learned by one of the following mechanisms:

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

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

- Related Documentation**
- [Understanding Static RP on page 181](#)
 - [Understanding RP Mapping with Anycast RP on page 185](#)
 - [Understanding the PIM Bootstrap Router on page 201](#)
 - [Understanding PIM Auto-RP on page 195](#)

Designated Router

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

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

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

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



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

Enabling PIM Sparse Mode

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

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

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

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

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

To configure the router properties for PIM sparse mode:

1. Configure the static RP router.

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

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

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

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

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

4. Monitor the operation of PIM sparse mode.

- [show pim interfaces](#)
- [show pim join](#)
- [show pim neighbors](#)
- [show pim rps](#)

Related Documentation

- [Understanding PIM Sparse Mode on page 143](#)

Configuring PIM Join Load Balancing

By default, PIM join messages are sent toward a source based on the RPF routing table check. If there is more than one equal-cost path toward the source, then one upstream interface is chosen to send the join message. This interface is also used for all downstream

traffic, so even though there are alternative interfaces available, the multicast load is concentrated on one upstream interface and routing device.

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

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

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

By default, when multiple PIM joins are received for different groups, all joins are sent to the same upstream gateway chosen by the unicast routing protocol. Even if there are multiple equal-cost paths available, these alternative paths are not utilized to distribute multicast traffic from the source to the various groups.

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

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

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

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

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

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

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

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

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

```
[edit protocols pim ]
user@host# set interface all mode sparse version 2
user@host# set join-load-balance
```

3. Then configure the static address of the RP.

```
[edit protocols pim rp]
user@host# set static address 10.10.10.1
```

4. Monitor the operation.

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

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

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

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

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

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

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

```
Interface: t1-0/2/3.0
```

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

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

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

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

- Related Documentation**
- [clear pim join-distribution](#)
 - [show pim interfaces on page 517](#)
 - [show pim neighbors on page 541](#)
 - [show pim source on page 552](#)

Modifying the Join State Timeout

This section describes how to configure the join state timeout.

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

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

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

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

The join timeout value can be from 210 through 420 seconds.

Related Documentation

- [join-prune-timeout on page 355](#)

Example: Enabling Join Suppression

This example describes how to enable PIM join suppression.

- [Requirements on page 151](#)
- [Overview on page 151](#)
- [Configuration on page 154](#)
- [Verification on page 155](#)

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See “[Enabling PIM Sparse Mode](#)” on [page 146](#).

Overview

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


This example includes the following statements:

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

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

[Figure 22 on page 153](#) shows the topology used in this example.

Figure 22: Join Suppression



The items in [Figure 22 on page 153](#) represent the following functions:

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

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

Configuration

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

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

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

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

1. Configure PIM sparse mode on the interfaces.

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

2. Enable the join suppression timer.

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

3. Configure the prune override interval value.

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

4. Configure the propagation delay of the link.

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

5. (Optional) Configure PIM tracing operations.

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

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

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

Results

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

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

Verification

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

- **show pim join extensive**

- [show multicast route extensive](#)

**Related
Documentation**

- [Example: Configuring the PIM Assert Timeout on page 222](#)
- [Example: Configuring PIM RPF Selection](#)
- [Example: Configuring the PIM SPT Threshold Policy on page 224](#)
- [Enabling PIM Sparse Mode on page 146](#)
- [PIM Overview on page 125](#)

CHAPTER 8

Using PIM Dense Mode and PIM Sparse-Dense Mode

- [Understanding PIM Dense Mode on page 157](#)
- [Understanding PIM Sparse-Dense Mode on page 159](#)
- [Mixing PIM Sparse and Dense Modes on page 159](#)
- [Configuring PIM Dense Mode Properties on page 160](#)
- [Configuring PIM Sparse-Dense Mode Properties on page 161](#)

Understanding PIM Dense Mode

PIM dense mode is less sophisticated than PIM sparse mode. PIM dense mode is useful for multicast LAN applications, the main environment for all dense mode protocols.


PIM dense mode implements the same flood-and-prune mechanism that DVMRP and other dense mode routing protocols employ. The main difference between DVMRP and PIM dense mode is that PIM dense mode introduces the concept of protocol independence. PIM dense mode can use the routing table populated by any underlying unicast routing protocol to perform reverse-path-forwarding (RPF) checks.

Internet service providers (ISPs) typically appreciate the ability to use any underlying unicast routing protocol with PIM dense mode because they do not need to introduce and manage a separate routing protocol just for RPF checks. While unicast routing protocols extended as multiprotocol BGP (MBGP) and Multitopology Routing in IS-IS (M-IS-IS) were later employed to build special tables to perform RPF checks, PIM dense mode does not require them.

PIM dense mode can use the unicast routing table populated by OSPF, IS-IS, BGP, and so on, or PIM dense mode can be configured to use a special multicast RPF table populated by MBGP or M-IS-IS when performing RPF checks.

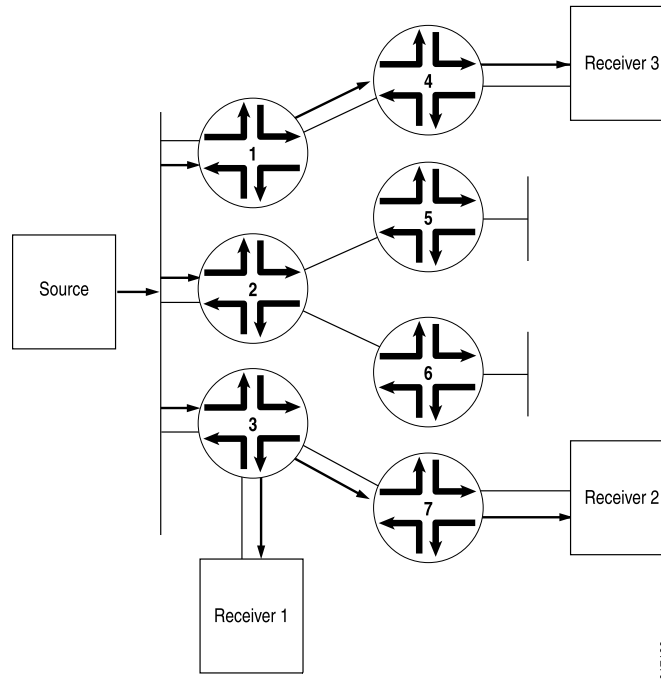
Unlike sparse mode, in which data is forwarded only to routing devices sending an explicit request, dense mode implements a *flood-and-prune* mechanism, similar to DVMRP. In PIM dense mode, there is no RP. A routing device receives the multicast data on the interface closest to the source, then forwards the traffic to all other interfaces (see [Figure 23 on page 158](#)).

Figure 23: Multicast Traffic Flooded from the Source Using PIM Dense Mode



Flooding occurs periodically. It is used to refresh state information, such as the source IP address and multicast group pair. If the routing device has no interested receivers for the data, and the OIL becomes empty, the routing device sends a prune message upstream to stop delivery of multicast traffic (see [Figure 24 on page 159](#)).

Figure 24: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic



Understanding PIM Sparse-Dense Mode

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

For information about PIM sparse-mode and PIM dense-mode rules, see [“Understanding PIM Sparse Mode” on page 143](#) and [“Understanding PIM Dense Mode” on page 157](#).

- Related Documentation**
- [Understanding PIM Sparse Mode on page 143](#)
 - [Understanding PIM Dense Mode on page 157](#)

Mixing PIM Sparse and Dense Modes

It is possible to mix PIM dense mode, PIM sparse mode, and PIM source-specific multicast (SSM) on the same network, the same routing device, and even the same interface. This is because modes are effectively tied to multicast groups, an IP multicast group address must be unique for a particular group's traffic, and scoping limits enforce the division between potential or actual overlaps.



NOTE: PIM sparse mode was capable of forming shortest-path trees (SPTs) already. Changes to PIM sparse mode to support PIM SSM mainly involved defining behavior in the SSM address range, because shared-tree behavior is prohibited for groups in the SSM address range.

A multicast routing device employing sparse-dense mode is a good example of mixing PIM modes on the same network or routing device or interface. Dense modes are easy to support because of the flooding, but scaling issues make dense modes inappropriate for Internet use beyond very restricted uses.

Configuring PIM Dense Mode Properties

In PIM dense mode (PIM-DM), the assumption is that almost all possible subnets have at least one receiver wanting to receive the multicast traffic from a source, so the network is flooded with traffic on all possible branches, then pruned back when branches do not express an interest in receiving the packets, explicitly (by message) or implicitly (time-out silence). LANs are appropriate networks for dense-mode operation.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default.

You can configure PIM dense mode globally or for a routing instance. This example shows how to configure the routing instance and how to specify that PIM dense mode use **inet.2** as its RPF routing table instead of **inet.0**.

To configure the router properties for PIM dense mode:

1. (Optional) Create an IPv4 routing table group so that interface routes are installed into two routing tables, **inet.0** and **inet.2**.

```
[edit routing-options rib-groups]
user@host# set pim-rg export-rib inet.0
user@host# set pim-rg import-rib [ inet.0 inet.2 ]
```

2. (Optional) Associate the routing table group with a PIM routing instance.

```
[edit routing-instances PIM.dense protocols pim]
user@host# set rib-group inet pim-rg
```

3. Configure the PIM interface. If you do not specify any interfaces, PIM is enabled on all router interfaces. Generally, you specify interface names only if you are disabling PIM on certain interfaces.

```
[edit routing-instances PIM.dense protocols pim]
user@host# set interface fe-0/0/1.0 mode dense
```



NOTE: You cannot configure both PIM and Distance Vector Multicast Routing Protocol (DVMRP) in forwarding mode on the same interface. You can configure PIM on the same interface only if you configured DVMRP in unicast-routing mode.

4. Monitor the operation of PIM dense mode by running the **show pim interfaces**, **show pim join**, **show pim neighbors**, and **show pim statistics** commands.

Related Documentation

- [Understanding PIM Dense Mode on page 157](#)
- [Example: Configuring a Dedicated PIM RPF Routing Table](#)

Configuring PIM Sparse-Dense Mode Properties

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 an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default.

You can configure PIM sparse-dense mode globally or for a routing instance. This example shows how to configure PIM sparse-dense mode globally on all interfaces, specifying that the groups 224.0.1.39 and 224.0.1.40 are using dense mode.

To configure the router properties for PIM sparse-dense mode:

1. Configure the dense-mode groups.

```
[protocols pim]
user@host# set dense-groups 224.0.1.39
user@host# set dense-groups 224.0.1.40
```

2. Configure all interfaces on the routing device to use sparse-dense mode. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

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

3. Monitor the operation of PIM sparse-dense mode by running the **show pim interfaces**, **show pim join**, **show pim neighbors**, and **show pim statistics** commands.

Related Documentation

- [Understanding PIM Sparse-Dense Mode on page 159](#)

CHAPTER 9

Using Source-Specific Multicast

- [Source-Specific Multicast Groups Overview on page 163](#)
- [Understanding PIM Source-Specific Mode on page 164](#)
- [PIM SSM on page 165](#)
- [Example: Configuring PIM SSM on a Network on page 167](#)
- [Example: Configuring an SSM-Only Domain on page 169](#)
- [Example: Configuring SSM Mapping on page 169](#)
- [Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 172](#)
- [Example: Configuring SSM Maps for Different Groups to Different Sources on page 175](#)

Source-Specific Multicast Groups Overview

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

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

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

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

SSM is typically supported with a subset of IGMPv3 and PIM sparse mode known as *PIM SSM*. Using SSM, a client can receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the client and the source, but builds the SPT without the help of an RP.

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

Understanding PIM Source-Specific Mode

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

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

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

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

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

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

[Table 8 on page 165](#).

Table 8: ASM and SSM Terminology

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

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

PIM SSM

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

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

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


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

Deploying SSM is easy. You need to configure PIM sparse mode on all router interfaces and issue the necessary SSM commands, including specifying IGMPv3 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group member

interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3, are used in PIM SSM. As sources become active and start sending multicast packets, interested receivers in the SSM group receive the multicast packets.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3), announcing a desire to join group G and source S (see [Figure 25 on page 166](#)). The directly connected PIM sparse-mode router, the receiver's DR, sends an (S,G) join message to its RPF neighbor for the source. Notice in [Figure 25 on page 166](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 25: Receiver Announces Desire to Join Group G and Source S




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 26 on page 166](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

Figure 26: Router 3 (Last-Hop Router) Joins the Source Tree



Using the source tree, multicast traffic is delivered to the subscribing host (see [Figure 27 on page 166](#)).

Figure 27: (S,G) State Is Built Between the Source and the Receiver




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

- Related Documentation**
- [Source-Specific Multicast Groups Overview on page 163](#)
 - [Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 172](#)

Example: Configuring PIM SSM on a Network

The following example shows how PIM SSM is configured between a receiver and a source in the network illustrated in [Figure 28 on page 167](#).

Figure 28: Network on Which to Configure PIM SSM



This example shows how to configure the IGMP version to IGMPv3 on all receiving host interfaces.

1. Enable IGMPv3 on all host-facing interfaces, and disable IGMP on the **fxp0.0** interface on Router 1.

```

user@router1# set protocols igmp interface all version 3
user@router1# set protocols igmp interface fxp0.0 disable

```



NOTE: When you configure IGMPv3 on a router, hosts on interfaces configured with IGMPv2 cannot join the source tree.

2. After the configuration is committed, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```

user@router1> show configuration protocol igmp

[edit protocols igmp]
interface all {
  version 3;
}
interface fxp0.0 {
  disable;
}

```

3. Use the **show igmp interface** command to verify that IGMP interfaces are configured.

```

user@router1> show igmp interface

```

Interface	State	Querier	Timeout	Version	Groups
fe-0/0/0.0	Up	198.58.3.245	213	3	0
fe-0/0/1.0	Up	198.58.3.241	220	3	0
fe-0/0/2.0	Up	198.58.3.237	218	3	0

```

Configured Parameters:
IGMP Query Interval (1/10 secs): 1250

```

```

IGMP Query Response Interval (1/10 secs): 100
IGMP Last Member Query Interval (1/10 secs): 10
IGMP Robustness Count: 2
Derived Parameters:
IGMP Membership Timeout (1/10 secs): 2600
IGMP Other Querier Present Timeout (1/10 secs): 2550

```

4. Use the **show pim join extensive** command to verify the PIM join state on Router 2 and Router 3 (the upstream routers).

```

user@router2> show pim join extensive
232.1.1.1      10.4.1.2      sparse
Upstream interface: fe-1/1/3.0
Upstream State: Local Source
Keepalive timeout: 209
Downstream Neighbors:
Interface: so-1/0/2.0
10.10.71.1     State: Join   Flags: S     Timeout: 209

```

5. Use the **show pim join extensive** command to verify the PIM join state on Router 1 (the router connected to the receiver).

```

user@router1> show pim join extensive
232.1.1.1      10.4.1.2      sparse
Upstream interface: so-1/0/2.0
Upstream State: Join to Source
Keepalive timeout: 209
Downstream Neighbors:
Interface: fe-0/2/3.0
10.3.1.1       State: Join   Flags: S     Timeout: Infinity

```



NOTE: IP version 6 (IPv6) multicast routers use the Multicast Listener Discovery (MLD) Protocol to manage the membership of hosts and routers in multicast groups and to learn which groups have interested listeners for each attached physical networks. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol. MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

Related Documentation

- [Example: Configuring SSM Mapping on page 169](#)

Example: Configuring an SSM-Only Domain

Deploying an SSM-only domain is much simpler than deploying an ASM domain because it only requires a few configuration steps. Enable PIM sparse mode on all interfaces by adding the **mode** statement at the **[edit protocols pim interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface. Then configure IGMPv3 on all host-facing interfaces by adding the **version** statement at the **[edit protocols igmp interface *interface-name*]** hierarchy level.

In the following example, the host-facing interface is **fe-0/1/2**:

```
[edit]
protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
  igmp {
    interface fe-0/1/2 {
      version 3;
    }
  }
}
```

Example: Configuring SSM Mapping

SSM mapping does not require that all hosts support IGMPv3. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report. This enables hosts running IGMPv1 or IGMPv2 to participate in SSM until the hosts transition to IGMPv3.

SSM mapping applies to all group addresses that match the policy, not just those that conform to SSM addressing conventions (232/8 for IPv4, ff30::/32 through ff3F::/32 for IPv6).

We recommend separate SSM maps for IPv4 and IPv6 if both address families require SSM support. If you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv4 context (using IGMP), only the IPv4 addresses in the list are used. If there are no such addresses, no action is taken. Similarly, if you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv6 context (using MLD), only the IPv6 addresses in the list are used. If there are no such addresses, no action is taken.

In this example, you create a policy to match the group addresses that you want to translate to IGMPv3. Then you define the SSM map that associates the policy with the source addresses where these group addresses are found. Finally, you apply the SSM map to one or more IGMP (for IPv4) or MLD (for IPv6) interfaces.

1. Create an SSM policy named **ssm-policy-example**. The policy terms match the IPv4 SSM group address 232.1.1.1/32 and the IPv6 SSM group address ff35::1/128. All other addresses are rejected.

```
user@router1# set policy-options policy-statement ssm-policy-example term A from
route-filter 232.1.1.1/32 exact
user@router1# set policy-options policy-statement ssm-policy-example term A then
accept
user@router1# set policy-options policy-statement ssm-policy-example term B from
route-filter ff35::1/128 exact
user@router1# set policy-options policy-statement ssm-policy-example term B then
accept
```

2. After the configuration is committed, use the **show configuration policy-options** command to verify the policy configuration.

```
user@host> show configuration policy-options

[edit policy-options]
policy-statement ssm-policy-example {
  term A {
    from {
      route-filter 232.1.1.1/32 exact;
    }
    then accept;
  }
  term B {
    from {
      route-filter ff35::1/128 exact;
    }
    then accept;
  }
  then reject;
}
```

The group addresses must match the configured policy for SSM mapping to occur.

3. Define two SSM maps, one called **ssm-map-ipv6-example** and one called **ssm-map-ipv4-example**, by applying the policy and configuring the source addresses as a multicast routing option.

```
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example policy
ssm-policy-example
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example source
fec0::1 fec0::12
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example policy
ssm-policy-example
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
10.10.10.4
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
192.168.43.66
```

4. After the configuration is committed, use the **show configuration routing-options** command to verify the policy configuration.

```
user@host> show configuration routing-options

[edit routing-options]
multicast {
  ssm-map ssm-map-ipv6-example {
    policy ssm-policy-example;
    source [ fec0::1 fec0::12 ];
  }
  ssm-map ssm-map-ipv4-example {
    policy ssm-policy-example;
    source [ 10.10.10.4 192.168.43.66 ];
  }
}
```

We recommend separate SSM maps for IPv4 and IPv6.

5. Apply SSM maps for IPv4-to-IGMP interfaces and SSM maps for IPv6-to-MLD interfaces:

```
user@host# set protocols igmp interface fe-0/1/0.0 ssm-map ssm-map-ipv4-example
user@host# set protocols mld interface fe-0/1/1.0 ssm-map ssm-map-ipv6-example
```

6. After the configuration is committed, use the **show configuration protocol** command to verify the IGMP and MLD protocol configuration.

```
user@router1> show configuration protocol

[edit protocols]
igmp {
  interface fe-0/1/0.0 {
    ssm-map ssm-map-ipv4-example;
  }
}
mld {
  interface fe-0/1/1.0 {
    ssm-map ssm-map-ipv6-example;
  }
}
```

7. Use the **show igmp interface** and the **show mld interface** commands to verify that the SSM maps are applied to the interfaces.

```
user@host> show igmp interface fe-0/1/0.0
Interface: fe-0/1/0.0
  Querier: 192.168.224.28
  State:      Up Timeout:      None Version: 2 Groups: 2
  SSM Map: ssm-map-ipv4-example

user@host> show mld interface fe-0/1/1.0
Interface: fe-0/1/1.0
  Querier: fec0:0:0:0:1::12
  State:      Up Timeout:      None Version: 2 Groups: 2
  SSM Map: ssm-map-ipv6-example
```

Example: Configuring Source-Specific Multicast Groups with Any-Source Override

This example shows how to extend source-specific multicast (SSM) group operations beyond the default IP address range of 232.0.0.0 through 232.255.255.255. This example also shows how to accept any-source multicast (ASM) join messages (*G) for group addresses that are within the default or configured range of SSM groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

- [Requirements on page 172](#)
- [Overview on page 172](#)
- [Configuration on page 173](#)
- [Verification on page 175](#)

Requirements


Before you begin, configure the router interfaces.

Overview

To deploy SSM, configure PIM sparse mode on all routing device interfaces and issue the necessary SSM commands, including specifying IGMPv3 or MLDv2 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group members interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3 and MLDv2, are used in PIM SSM. Only sources that are specified send traffic to the SSM group.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3 or MLDv2) to join group G and source S (see [Figure 29 on page 172](#)). The directly connected PIM sparse-mode router, the receiver's designated router (DR), sends an (S,G) join message to its reverse-path forwarding (RPF) neighbor for the source. Notice in [Figure 29 on page 172](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 29: Receiver Sends Messages to Join Group G and Source S




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 30 on page 173](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

Figure 30: Router 3 (Last-Hop Router) Joins the Source Tree



Using the source tree, multicast traffic is delivered to the subscribing host (see [Figure 31 on page 173](#)).


Figure 31: (S,G) State Is Built Between the Source and the Receiver



SSM can operate in include mode or in exclude mode. In exclude mode the receiver specifies a list of sources that it does not want to receive the multicast group traffic from. The routing device forwards traffic to the receiver from any source except the sources specified in the exclusion list. The receiver accepts traffic from any sources except the sources specified in the exclusion list.

This example works with the simple RPF topology shown in [Figure 32 on page 173](#).

Figure 32: Simple RPF Topology



Configuration

CLI Quick Configuration

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

```
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 10.255.72.46
set protocols pim rp local group-ranges 239.0.0.0/24
set protocols pim interface fe-1/0/0.0 mode sparse
set protocols pim interface lo0.0 mode sparse
set routing-options multicast ssm-groups 232.0.0.0/8
```

```
set routing-options multicast ssm-groups 239.0.0.0/8
set routing-options multicast asm-override-ssm
```

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

To configure an RPF policy:

1. Configure OSPF.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface fxp0.0 disable
user@host# set area 0.0.0.0 interface all
```

2. Configure PIM sparse mode.

```
[edit protocols pim]
user@host# set rp local address 10.255.72.46
user@host# set rp local group-ranges 239.0.0.0/24
user@host# set interface fe-1/0/0.0 mode sparse
user@host# set interface lo0.0 mode sparse
```

3. Configure additional SSM groups.

```
[edit routing-options]
user@host# set ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ]
```

4. Configure the RP to accept ASM join messages for groups within the SSM address range.

```
[edit routing-options]
user@host# set multicast asm-override-ssm
```

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

```
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols** and **show routing-options** commands.

```
user@host# show protocols
ospf {
  area 0.0.0.0 {
    interface fxp0.0 {
      disable;
    }
    interface all;
  }
}
pim {
  rp {
    local {
      address 10.255.72.46;
      group-ranges {
```

```

        239.0.0.0/24;
    }
}
interface fe-1/0/0.0 {
    mode sparse;
}
interface lo0.0 {
    mode sparse;
}
}

user@host# show routing-options
multicast {
    ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ];
    asm-override-ssm;
}

```

Verification

To verify the configuration, run the following commands:

- [show igmp group](#)
- [show igmp statistics](#)
- [show pim join](#)

Related Documentation

- [Source-Specific Multicast Groups Overview on page 163](#)

Example: Configuring SSM Maps for Different Groups to Different Sources

- [Multiple SSM Maps and Groups for Interfaces on page 175](#)
- [Example: Configuring Multiple SSM Maps Per Interface on page 175](#)

Multiple SSM Maps and Groups for Interfaces

You can configure multiple source-specific multicast (SSM) maps so that different groups map to different sources, which enables a single multicast group to map to different sources for different interfaces.

Example: Configuring Multiple SSM Maps Per Interface

This example shows how to assign more than one SSM map to an IGMP interface.

- [Requirements on page 176](#)
- [Overview on page 176](#)
- [Configuration on page 176](#)
- [Verification on page 178](#)

Requirements

This example requires Junos OS Release 11.4 or later.

Overview

In this example, you configure a routing policy, POLICY-ipv4-example1, that maps multicast group join messages over an IGMP logical interface to IPv4 multicast source addresses based on destination IP address as follows:

Routing Policy Name	Multicast Group Join Messages for a Route Filter at This Destination Address	Multicast Source Addresses
POLICY-ipv4-example1 term 1	232.1.1.1	10.10.10.4, 192.168.43.66
POLICY-ipv4-example1 term 2	232.1.1.2	10.10.10.5, 192.168.43.67

You apply routing policy POLICY-ipv4-example1 to IGMP logical interface fe-0/1/0.0.

Configuration

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

To configure this example, perform the following task:

CLI Quick Configuration

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

```

set policy-options policy-statement POLICY-ipv4-example1 term 1 from route-filter
  232.1.1.1/32 exact
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source
  10.10.10.4
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source
  192.168.43.66
set policy-options policy-statement POLICY-ipv4-example1 term 1 then accept
set policy-options policy-statement POLICY-ipv4-example1 term 2 from route-filter
  232.1.1.2/32 exact
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source
  10.10.10.5
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source
  192.168.43.67
set policy-options policy-statement POLICY-ipv4-example1 term 2 then accept
set protocols igmp interface fe-0/1/0.0 ssm-map-policy POLICY-ipv4-example1

```

Step-by-Step Procedure

To configure multiple SSM maps per interface:

1. Configure protocol-independent routing options for route filter 232.1.1.1, and specify the multicast source addresses to which matching multicast groups are to be mapped.

```
[edit policy-options policy-statement POLICY-ipv4-example1 term 1]
user@host# set from route-filter 232.1.1.1/32 exact
user@host# set then ssm-source 10.10.10.4
user@host# set then ssm-source 192.168.43.66
user@host# set then accept
```

2. Configure protocol-independent routing options for route filter 232.1.1.2, and specify the multicast source addresses to which matching multicast groups are to be mapped.

```
[edit policy-options policy-statement POLICY-ipv4-example1 term 2]
user@host# set from route-filter 232.1.1.2/32 exact
user@host# set then ssm-source 10.10.10.5
user@host# set then ssm-source 192.168.43.67
user@host# set then accept
```

3. Apply the policy map POLICY-ipv4-example1 to IGMP logical interface fe-0/1/1/0.

```
[edit protocols igmp interface fe-0/1/0.0]
user@host# set ssm-map-policy POLICY-ipv4-example1
```

Results After the configuration is committed, confirm the configuration by entering the **show policy-options** and **show protocols** configuration mode commands. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
user@host# show policy-options
policy-statement POLICY-ipv4-example1 {
  term 1 {
    from {
      route-filter 232.1.1.1/32 exact;
    }
    then {
      ssm-source [ 10.10.10.4 192.168.43.66 ];
      accept;
    }
  }
  term 2 {
    from {
      route-filter 232.1.1.2/32 exact;
    }
    then {
      ssm-source [ 10.10.10.5 192.168.43.67 ];
      accept;
    }
  }
}

user@host# show protocols
igmp {
  interface fe-0/1/0.0 {
```

```
        ssm-map-policy POLICY-ipv4-example1;  
    }  
}
```

Verification

Confirm that the configuration is working properly.

- [Displaying Information About IGMP-Enabled Interfaces on page 178](#)
- [Displaying the PIM Groups on page 178](#)
- [Displaying the Entries in the IP Multicast Forwarding Table on page 178](#)

Displaying Information About IGMP-Enabled Interfaces

Purpose Verify that the SSM map policy POLICY-ipv4-example1 is applied to logical interface fe-0/1/0.0.

Action Use the [show igmp interface](#) operational mode command for the IGMP logical interface to which you applied the SSM map policy.

```
user@host> show igmp interface  
Interface: fe-0/1/0.0  
  Querier: 10.111.30.1  
  State:      Up Timeout:   None Version:  2 Groups:      2  
  SSM Map Policy: POLICY-ipv4-example1;
```

```
Configured Parameters:  
IGMP Query Interval: 125.0  
IGMP Query Response Interval: 10.0  
IGMP Last Member Query Interval: 1.0  
IGMP Robustness Count: 2
```

```
Derived Parameters:  
IGMP Membership Timeout: 260.0  
IGMP Other Querier Present Timeout: 255.0
```

The command output displays the name of the IGMP logical interface (fe-0/1/0.0), which is the address of the routing device that has been elected to send membership queries and group information.

Displaying the PIM Groups

Purpose Verify the Protocol Independent Multicast (PIM) source and group pair (S,G) entries.

Action Use the [show pim join extensive 232.1.1.1](#) operational mode command to display the PIM source and group pair (S,G) entries for the 232.1.1.1 group.

Displaying the Entries in the IP Multicast Forwarding Table

Purpose Verify that the IP multicast forwarding table displays the multicast route state.

Action Use the [show multicast route extensive](#) operational mode command to display the entries in the IP multicast forwarding table to verify that the **Route state** is active and that the **Forwarding state** is forwarding.

- Related Documentation**
- *Example: Configuring Source-Specific Multicast*
 - *Example: Configuring Source-Specific Draft-Rosen 7 Multicast VPNs*

CHAPTER 10

Using Static RP

- [Understanding Static RP on page 181](#)
- [Configuring Local PIM RPs on page 181](#)
- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 183](#)

Understanding Static RP

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

A static configuration is simple and convenient. However, if the statically defined RP router becomes unreachable, there is no automatic failover to another RP router. To remedy this problem, you can use anycast RP.

Related Documentation

- [Configuring Local PIM RPs on page 181](#)
- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 183](#)

Configuring Local PIM RPs

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

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

To configure the routing device's RP properties:

1. Configure the routing instance as the local RP.

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

2. Configure the IP protocol family and IP address.

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

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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

- Related Documentation**
- [PIM Overview on page 125](#)
 - [Understanding MLD on page 75](#)

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

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

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

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

To configure the static RP address:

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

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



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

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

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

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



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

3. (Optional) Set the group address range.

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

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

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

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

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

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

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

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

- Related Documentation**
- [PIM Overview on page 125](#)
 - [Understanding MLD on page 75](#)

CHAPTER 11

Using Anycast RP

- [Understanding RP Mapping with Anycast RP on page 185](#)
- [Example: Configuring PIM Anycast With or Without MSDP on page 186](#)
- [Configuring a PIM Anycast RP Router with MSDP on page 189](#)
- [Configuring a PIM Anycast RP Router Using Only PIM on page 190](#)
- [Configuring All PIM Anycast Non-RP Routers on page 191](#)
- [Example: Configuring Multiple RPs in a Domain with Anycast RP on page 192](#)

Understanding RP Mapping with Anycast RP

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

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

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

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

Related Documentation

- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 183](#)
- [Example: Configuring Multiple RPs in a Domain with Anycast RP on page 192](#)

- [Example: Configuring PIM Anycast With or Without MSDP on page 186](#)

Example: Configuring PIM Anycast With or Without MSDP

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

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

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

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

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

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

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

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

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

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

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

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

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

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



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

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

```
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
```

```

        family inet {
            address 198.58.3.254/32 {
                primary;
            }
            address 198.58.3.253/32;
        }
    }
}

```

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

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

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

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

```

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

```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at

the `[edit protocols pim rp static]` hierarchy level. Include the **version** statement at the `[edit protocols pim rp static address]` hierarchy level to specify PIM version 2.

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

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

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

Configuring a PIM Anycast RP Router with MSDP

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

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

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

```
        interface fxp0.0 {
            disable;
        }
    }
}
```

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

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

Configuring a PIM Anycast RP Router Using Only PIM

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

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

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

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

Use the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are

configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

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

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

Configuring All PIM Anycast Non-RP Routers

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

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

Example: Configuring Multiple RPs in a Domain with Anycast RP

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

- [Requirements on page 192](#)
- [Overview on page 192](#)
- [Configuration on page 192](#)
- [Verification on page 194](#)

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See [“Enabling PIM Sparse Mode” on page 146](#).

Overview

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

Configuration

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

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

Non-RP Routers	<pre>set protocols pim rp static address 10.1.1.2</pre>
-----------------------	---

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

To configure anycast RP:

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

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

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

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

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

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

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

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

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

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

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

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

```
user@host# commit
```

Results

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

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

```
        address 10.1.1.2/32;
    }
}
}
```

On the RP routers:

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

On the non-RP routers:

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

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

Verification

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

Related Documentation

- [Example: Configuring PIM Anycast With or Without MSDP on page 186](#)
- [Understanding PIM Sparse Mode on page 143](#)
- [Understanding RP Mapping with Anycast RP on page 185](#)

CHAPTER 12

Using Auto-RP

- [Understanding PIM Auto-RP on page 195](#)
- [Configuring PIM Auto-RP on page 195](#)

Understanding PIM Auto-RP

You can configure a more dynamic way of assigning rendezvous points (RPs) in a multicast network by means of auto-RP. When you configure auto-RP for a router, the router learns the address of the RP in the network automatically and has the added advantage of operating in PIM version 1 and version 2.

Although auto-RP is a nonstandard (non-RFC-based) function that typically uses dense mode PIM to advertise control traffic, it provides an important failover advantage that simple static RP assignment does not. You can configure multiple routers as RP candidates. If the elected RP fails, one of the other preconfigured routers takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

Related Documentation

- [Configuring PIM Auto-RP on page 195](#)

Configuring PIM Auto-RP

For correct operation, every multicast router within a PIM domain must be able to map a particular multicast group address to the same rendezvous point (RP). The auto-RP mechanism is one way that a multicast router can learn the set of group-to-RP mappings. Auto-RP automatically distributes mapping information to routing devices. It simplifies use of multiple RPs for different multicast group ranges, thus allowing multiple RPs to act as backups for each other. Auto-RP relies on a router to act as the RP mapping agent. Potential RPs announce themselves to the mapping agent, and the mapping agent resolves any conflicts.

The mapping agent sends the multicast group-RP mapping information to the other routers using PIM dense mode. The specific groups used are 224.0.1.39 and .40. The first (.39) is used to advertise, the second (.40) is used for discovery. Because PIM dense mode is necessary to enable auto-RP to work, which in turns enables PIM sparse mode to work, you must configure PIM sparse-dense mode in the PIM domains that use auto-RP.

Although auto-RP is a nonstandard (non-RFC-based) function requiring dense mode PIM to advertise control traffic, it provides an important failover advantage that static

RP assignment does not. That is, you can configure multiple routing devices as RP candidates. If the elected RP fails, one of the other preconfigured routing devices takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

Auto-RP operates in PIM version 1 and version 2.

In most cases, how the routing device handles auto-RP discovery, announce, or mapping messages depends on whether the routing device is an RP (configured as local RP) or not. [Table 9 on page 196](#) shows how the routing device behaves depending on the local RP configuration.

Table 9: Local RP and Auto-RP Message Types

Auto-RP Message Type	Local RP?	Routing Device Behavior
discovery	No	Listen for auto-RP mapping messages.
discovery	Yes	Listen for auto-RP mapping messages.
announce	No	Listen for auto-RP mapping messages.
announce	Yes	Listen for auto-RP mapping messages. Send auto-RP announce messages.
mapping	No	Listen for auto-RP mapping messages. Listen for auto-RP announce messages. If elected mapping agent, send auto-RP mapping messages.
mapping	Yes	Listen for auto-RP mapping messages. Send auto-RP announce messages. Listen for auto-RP announce messages. If elected mapping agent, send auto-RP mapping messages.



NOTE: If the routing device receives auto-RP announcements split across multiple messages, the routing device loses the information in the previous part of the message as soon as the next part of the message is received.

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

To configure auto-RP properties:

1. Configure PIM in sparse-dense mode on all routing devices in the PIM domain.

```
[edit protocols pim]
user@host# edit
user@host# set interface all mode sparse-dense
```


This configuration allows the routing device to operate in sparse mode for most groups and dense mode for others. The default is to operate in sparse mode unless the routing device is specifically informed of a dense mode group.

2. Configure a routable loopback interface address on all routing devices in the PIM domain.

The routing device joins the auto-RP groups on the configured interfaces and on the loopback interface **lo0.0**. For auto-RP to work correctly, configure a routable IP address on the loopback interface. You cannot use the loopback address 127.0.0.1. Also, you must enable PIM sparse-dense mode on the **lo0.0** interface if you do not specify **interface all**.

```
[edit interfaces lo0.0 unit 0 family inet]
user@host# set address 192.168.0.3 preferred
```

3. Configure the two multicast dense groups on all the routing devices.

Auto-RP requires multicast flooding to announce potential RP candidates and to discover the elected RPs in the network. Multicast flooding occurs through a PIM dense mode model, where group 224.0.1.39 is used for **announce** messages and group 224.0.1.40 is used for **discovery** messages.

```
[edit protocols pim]
user@host# set dense-groups 224.0.1.39/32
user@host# set dense-groups 224.0.1.40/32
```



TIP: Step 3 is required. When auto-RP is enabled, the auto-RP announce group (224.0.1.39) and auto-RP-discovery group (224.0.1.40) must be configured explicitly as dense groups. When the auto-RP discovery group is not configured as a dense group, auto-RP is not enabled. When the auto-RP announce group is not configured as a dense group, auto-RP is enabled in the discovery mode only, and mapping and announce modes are disabled.

4. Configure the auto-RP **announce** option.

At least one routing device in the PIM domain must announce auto-RP messages and at least one must map them, or you can configure a routing device to perform both functions.

When a routing device sends announce messages in the network, it is advertising itself as a candidate RP. A routing device configured with this option must also be configured as an RP, or announce messages are not sent.

```
[edit protocols pim rp]
user@host# set local address 192.168.0.1
user@host# set auto-rp announce
```



NOTE: You cannot include the **auto-rp announce** option at the **[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim]** hierarchy level.

5. Configure the auto-RP mapping agent.

The mapping agent sends discovery messages to the network, informing all routing devices in a multicast group of which RP to use. If the mapping agent is also an RP, the **mapping** option also allows the routing device to send auto-RP announcements (mapping on an RP allows the routing device to perform both the announcement and mapping functions).

```
[edit protocols pim rp]
user@host# set auto-rp mapping
```

If the mapping agent is also an RP, configure the mapping agent as a local RP.

```
[edit protocols pim rp]
user@host# set local address 192.168.0.2
```

6. Configure mapping agent election.

If you configure the **mapping** option on more than one routing device in the PIM domain, configure mapping agent election on each potential mapping agent.

Auto-RP specifications state that mapping agents do not send mapping messages if they receive messages from a mapping agent with a higher IP address. However, some vendors' mapping agents continue to announce mappings, even in the presence of higher-addressed mapping agents. In other words, some mapping agents will always send mapping messages.

The default auto-RP operation is to perform mapping agent election. To explicitly configure mapping agent election, you can include the **mapping-agent-election** statement. When this option is configured, the mapping agent will stop sending mapping messages if it receives messages from a mapping agent with a higher IP address.

```
[edit protocols pim rp]
user@host# set auto-rp mapping mapping-agent-election
```

Mapping message suppression is disabled with the **no-mapping-agent-election** statement. When this option is configured, the mapping agent will always send mapping messages even in the presence of higher-addressed mapping agents.

To disable mapping agent election for compatibility with other vendors' equipment, include the **no-mapping-agent-election** statement.

```
[edit protocols pim rp]
user@host# set auto-rp mapping no-mapping-agent-election
```

7. Configure the remaining routing devices in the PIM domain to discover the RP.

Discovery enables the routing devices to receive and process discovery messages from the mapping agent. This is the most basic auto-RP option.

```
[edit protocols pim rp]
user@host# set auto-rp discovery
```

8. Monitor the operation of PIM auto-RP routers by running the following commands:

- **show pim interfaces**
- **show pim rps**

- [show pim rps](#)

9. Issue the **show pim rps extensive** command to see information about how an RP is learned, what groups it handles, and the number of groups actively using the RP.

```

user@host> show pim rps extensive
RP: 192.168.5.1
Learned from 192.168.5.1 via: auto-rp
Time Active: 00:34:29
Holdtime: 150 with 108 remaining
Device Index: 6
Subunit: 32769
Interface: pd-0/0/0.32769
Group Ranges:
    224.0.0.0/4
Active groups using RP:
    224.2.2.100
    total 1 groups active
Register State for RP:
Group      Source FirstHop      RP Address      StateRP address Type Holdtime
Timeout

```

In the example, the RP at 192.168.5.1 was learned through auto-RP. The RP is able to support all groups in the 224.0.0.0/4 range (all possible groups). The local router has sent PIM control traffic for the 224.2.2.100 group to the RP.

Additionally, the presence of a Tunnel Physical Interface Card (PIC) in an RP router creates a de-encapsulation interface, which allows the RP to receive multicast traffic from the source. This interface is indicated by **pd-0/0/0.32769**.

Related Documentation

- [Understanding PIM Sparse Mode on page 143](#)
- [show pim interfaces on page 517](#)
- [show pim rps on page 545](#)

Using PIM Bootstrap Router

- [Understanding the PIM Bootstrap Router on page 201](#)
- [Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 201](#)
- [Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 203](#)
- [Example: Configuring PIM BSR Filters on page 203](#)

Understanding the PIM Bootstrap Router

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

**Related
Documentation**

- [Configuring PIM Bootstrap Properties for IPv4 or IPv6](#)

Configuring PIM Bootstrap Properties for IPv4 or IPv6

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

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

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

To configure the bootstrap router properties:

1. Configure the bootstrap priority.

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



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

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

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

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

3. Configure the policies.

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

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

Related Documentation

- [Understanding PIM Sparse Mode on page 143](#)

- [Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 203](#)
- [show pim bootstrap on page 515](#) in the CLI Explorer

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

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

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

Example: Configuring PIM BSR Filters

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

```
protocols {
  pim {
    rp {
      bootstrap-import no-bsr;
      bootstrap-export no-bsr;
    }
  }
}
```

```
policy-options {  
  policy-statement no-bsr {  
    then reject;  
  }  
}
```


Using PIM Filtering

- [Understanding Multicast Message Filters on page 205](#)
- [Filtering MAC Addresses on page 206](#)
- [Filtering RP and DR Register Messages on page 206](#)
- [Configuring Interface-Level PIM Neighbor Policies on page 207](#)
- [Filtering Outgoing PIM Join Messages on page 208](#)
- [Filtering Incoming PIM Join Messages on page 209](#)
- [Configuring Register Message Filters on a PIM RP and DR on page 211](#)

Understanding Multicast Message Filters

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

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

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

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



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



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

**Related
Documentation**

- [Filtering MAC Addresses on page 206](#)
- [Filtering RP and DR Register Messages on page 206](#)
- [Filtering MSDP SA Messages on page 234](#)

Filtering MAC Addresses

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

No configuration is necessary for MAC filters.

Filtering RP and DR Register Messages

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

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

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

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

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

- Group multicast address
- Source address

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

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

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

Related Documentation

- [Understanding RP Mapping with Anycast RP on page 185](#)
- [Configuring Register Message Filters on a PIM RP and DR on page 211](#)

Configuring Interface-Level PIM Neighbor Policies

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

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

To configure a policy to filter unwanted PIM neighbors:

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

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

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

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

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

- Related Documentation**
- [Understanding PIM Sparse Mode on page 143](#)
 - [show pim statistics on page 555](#)

Filtering Outgoing PIM Join Messages

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

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

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

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

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

1. In configuration mode, create the policy.

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

```

user@host# set policy-options policy-statement block-groups term t1 then reject
user@host# set policy-options policy-statement block-groups term last then accept

```

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

```

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

```

3. Apply the PIM join and prune message filter.

```

user@host> set protocols pim export block-groups

```

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

```

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

Rx Joins/Prunes filtered    0

Tx Joins/Prunes filtered    254

```

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

Related Documentation • [Filtering Incoming PIM Join Messages on page 209](#)

Filtering Incoming PIM Join Messages

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

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

See [Table 10 on page 210](#) for a list of match conditions.

Table 10: PIM Join Filter Match Conditions

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

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

To filter incoming PIM join messages:

1. Configure the policy.

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

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

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

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

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

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

Related Documentation

- [Understanding Multicast Administrative Scoping](#)
- [Filtering Outgoing PIM Join Messages on page 208](#)
- [show pim join on page 520](#) in the [CLI Explorer](#)
- [show policy](#) in the [CLI Explorer](#)

Configuring Register Message Filters on a PIM RP and DR

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

Register messages have the following purposes:

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

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

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

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

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

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

1. On the RP, configure the policy.

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

2. Apply the policy to the RP.

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

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

1. On the DR, configure the policy.

```
[edit policy-options policy-statement outgoing-policy-for-rp]
```

```
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.10.1/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the DR.

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

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

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

1. On the RP, configure the policies.

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

[edit policy-options policy-statement accept_224_1_1_5]
user@host# set term one from route-filter 224.1.1.5/32 exact
user@host# set term one from source-address-filter 10.10.94.2/32 exact
user@host# set term one then accept
user@host# set term two then reject
user@host# exit
```

2. Apply the policies to the RP.

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

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

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

Related Documentation

- [PIM Sparse Mode Source Registration on page 215](#)
- [Filtering RP and DR Register Messages on page 206](#)
- [show pim statistics on page 555](#)

CHAPTER 15

Using PIM RPT and SPT Cutover

- [Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees on page 213](#)
- [Building an RPT Between the RP and Receivers on page 214](#)
- [PIM Sparse Mode Source Registration on page 215](#)
- [Multicast Shortest-Path Tree on page 218](#)
- [SPT Cutover on page 219](#)
- [SPT Cutover Control on page 222](#)
- [Example: Configuring the PIM Assert Timeout on page 222](#)
- [Example: Configuring the PIM SPT Threshold Policy on page 224](#)

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

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

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

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

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

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

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

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

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

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


Building an RPT Between the RP and Receivers

The RPT is the path between the RP and receivers (hosts) in a multicast group (see [Figure 33 on page 215](#)). The RPT is built by means of a PIM join message from a receiver's DR:

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

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

Figure 33: Building an RPT Between the RP and the Receiver




PIM Sparse Mode Source Registration

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

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


1. The source becomes active, sending out multicast packets on the LAN to which it is attached. The source's DR receives the packets and encapsulates them in a PIM register message, which it sends to the RP router (see [Figure 34 on page 216](#)).
2. When the RP router receives the PIM register message from the source, it sends a PIM join message back to the source.

Figure 34: PIM Register Message and PIM Join Message Exchanged




3. The source's DR receives the PIM join message and begins sending traffic down the SPT toward the RP router (see [Figure 35 on page 217](#)).
4. Once traffic is received by the RP router, it sends a register stop message to the source's DR to stop the register process.

Figure 35: Traffic Sent from the Source to the RP Router



- The RP router sends the multicast traffic down the RPT toward the receiver (see [Figure 36 on page 217](#)).

Figure 36: Traffic Sent from the RP Router Toward the Receiver



Multicast Shortest-Path Tree

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

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

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

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

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

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

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


SPTs are always the shortest path, but they are not necessarily short. That is, sources and receivers tend to be on the periphery of a router network, not on the backbone, and multicast distribution trees have a tendency to sprawl across almost every router in the network. Because multicast traffic can overwhelm a slow interface, and one packet can easily become a hundred or a thousand on the opposite side of the backbone, it makes sense to provide a shared tree as a distribution tree so that the multicast source can be located more centrally in the network, on the backbone. This sharing of distribution trees with roots in the core network is accomplished by a multicast rendezvous point. For more information about RPs, see [“Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees” on page 213](#).

SPT Cutover

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


1. Once the receiver's DR receives the first multicast packet from the source, the DR sends a PIM join message to its RPF neighbor (see [Figure 37 on page 219](#)).
2. The source's DR receives the PIM join message, and an additional (S,G) state is created to form the SPT.
3. Multicast packets from that particular source begin coming from the source's DR and flowing down the new SPT to the receiver's DR. The receiver's DR is now receiving two copies of each multicast packet sent by the source—one from the RPT and one from the new SPT.

Figure 37: Receiver DR Sends a PIM Join Message to the Source




4. To stop duplicate multicast packets, the receiver's DR sends a PIM prune message toward the RP router, letting it know that the multicast packets from this particular source coming in from the RPT are no longer needed (see [Figure 38 on page 220](#)).

Figure 38: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router




5. The PIM prune message is received by the RP router, and it stops sending multicast packets down to the receiver's DR. The receiver's DR is getting multicast packets only for this particular source over the new SPT. However, multicast packets from the source are still arriving from the source's DR toward the RP router (see [Figure 39 on page 220](#)).

Figure 39: RP Router Receives PIM Prune Message




6. To stop the unneeded multicast packets from this particular source, the RP router sends a PIM prune message to the source's DR (see [Figure 40 on page 221](#)).

Figure 40: RP Router Sends a PIM Prune Message to the Source DR



7. The receiver's DR now receives multicast packets only for the particular source from the SPT (see [Figure 41 on page 221](#)).

Figure 41: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router



SPT Cutover Control

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

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

Example: Configuring the PIM Assert Timeout

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

- [Requirements on page 222](#)
- [Overview on page 222](#)
- [Configuration on page 224](#)

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See “[Enabling PIM Sparse Mode](#)” on [page 146](#).

Overview

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

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

In [Figure 42 on page 223](#), both routing devices R1 and R2 forward multicast packets for the same (S,G) entry on a network. Both devices detect this situation and both devices send assert messages on the Ethernet network. An assert message contains, in addition to a source address and group address, a unicast cost metric for sending packets to the

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


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

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

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

Figure 42 on page 223 shows the topology for this example.

Figure 42: PIM Assert Topology



9040614

Configuration

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

To configure an assert timeout:

1. Configure the timeout period, in seconds.

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

2. (Optional) Trace assert messages.

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

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

```
user@host# commit
```

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

- `show pim join`
- `show pim statistics`

Related Documentation

- [Configuring PIM Trace Options on page 130](#)
- [SPT Cutover on page 219](#)
- [SPT Cutover Control on page 222](#)

Example: Configuring the PIM SPT Threshold Policy

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

- [Requirements on page 224](#)
- [Overview on page 225](#)
- [Configuration on page 226](#)
- [Verification on page 228](#)

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.

- Configure PIM Sparse Mode on the interfaces. See [“Enabling PIM Sparse Mode” on page 146](#).

Overview

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

This example includes the following settings:

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

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

Note these points when configuring the SPT threshold policy:

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

Note these points when configuring the SPT threshold policy:

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

Note these points when configuring the SPT threshold policy:

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

Configuration

CLI Quick Configuration

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

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

Step-by-Step Procedure

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

To configure an SPT threshold policy:

1. Apply the policy.

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

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

2. Configure the policy.

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

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

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

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

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

Results

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

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

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

Verification

To verify the configuration, run the [show pim join](#) command.

Related Documentation

- [SPT Cutover Control on page 222](#)

PART 3

Configuring MSDP

- [Using MSDP on page 231](#)

CHAPTER 16

Using MSDP

- [Understanding MSDP on page 231](#)
- [Configuring MSDP on page 232](#)
- [Filtering MSDP SA Messages on page 234](#)
- [Tracing MSDP Protocol Traffic on page 234](#)
- [Configuring the Interface to Accept Traffic from a Remote Source on page 236](#)
- [Example: Configuring MSDP on page 237](#)
- [Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 238](#)
- [Example: Configuring PIM Anycast With or Without MSDP on page 244](#)
- [Configuring a PIM Anycast RP Router with MSDP on page 247](#)

Understanding MSDP

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

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

The MSDP peer-RPF check is different from the normal RPF checks done by non-MSDP multicast routers. The goal of the peer-RPF check is to stop source-active messages from looping. Router R accepts source-active messages originated by Router S only from neighbor Router N or an MSDP mesh group member. For more information about configuring MSDP mesh groups, see [“Example: Configuring MSDP with Active Source Limits and Mesh Groups” on page 238](#).

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

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

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

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

Related Documentation

- [Configuring MSDP on page 232](#)

Configuring MSDP

To configure the Multicast Source Discovery Protocol (MSDP), include the **msdp** statement:

```
msdp {  
  disable;  
  active-source-limit {  
    maximum number;  
    threshold number;  
  }  
  data-encapsulation (disable | enable);  
  export [ policy-names ];  
  group group-name {  
    ... group-configuration ...  
  }  
}
```

```

hold-time seconds;
import [ policy-names ];
local-address address;
keep-alive seconds;
peer address {
    ... peer-configuration ...
}
rib-group group-name;
source ip-prefix </prefix-length> {
    active-source-limit {
        maximum number;
        threshold number;
    }
}
sa-hold-time seconds;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
group group-name {
    disable;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    mode (mesh-group | standard);
    peer address {
        ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
        just following ...
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
peer address {
    disable;
    active-source-limit {
        maximum number;
        threshold number;
    }
    authentication-key peer-key;
    default-peer;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
}

```

You can include this statement at the following hierarchy levels:

- [\[edit protocols\]](#)
- [\[edit routing-instances *routing-instance-name* protocols\]](#)
- [\[edit logical-systems *logical-system-name* protocols\]](#)
- [\[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols\]](#)

By default, MSDP is disabled.

- Related Documentation**
- [Example: Configuring MSDP in a Routing Instance](#)
 - [Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 238](#)

Filtering MSDP SA Messages

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

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



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

- Related Documentation**
- [Filtering Incoming PIM Join Messages on page 209](#)
 - [Example: Configuring PIM BSR Filters on page 203](#)

Tracing MSDP Protocol Traffic

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

Flag	Description
all	Trace all operations.

Flag	Description
general	Trace general events.
keepalive	Trace keepalive messages.
normal	Trace normal events.
packets	Trace all MSDP packets.
policy	Trace policy processing.
route	Trace MSDP changes to the routing table.
source-active	Trace source-active packets.
source-active-request	Trace source-active request packets.
source-active-response	Trace source-active response packets.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

You can configure MSDP tracing for all peers, for all peers in a particular group, or for a particular peer.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on MSDP peers in a particular group. To configure tracing operations for MSDP:

1. (Optional) Configure tracing by including the **traceoptions** statement at the **[edit routing-options]** hierarchy level and set the **all-packets-trace** and **all** flags to trace all protocol packets.

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

2. Configure the filename for the MSDP trace file.

```
[edit protocols msdp group groupa traceoptions]
user@host# set file msdp-trace
```

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

```
[edit protocols msdp group groupa traceoptions]
user@host# set file files 5
```

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

```
[edit protocols msdp group groupa traceoptions]
```

```
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols msdp group groupa traceoptions]
```

```
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with the source-active cache for **groupa**. The following example shows how to trace messages associated with the group address.

```
[edit protocols msdp group groupa traceoptions]
```

```
user@host# set flag source-active | match 230.0.0.3
```

7. View the trace file.

```
user@host> file list /var/log
```

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


Related Documentation

- [Understanding MSDP on page 231](#)
- *Tracing and Logging Junos OS Operations in the Junos OS Administration Library for Routing Devices*

Configuring the Interface to Accept Traffic from a Remote Source

You can configure an incoming interface to accept multicast traffic from a remote source. A remote source is a source that is not on the same subnet as the incoming interface. [Figure 43 on page 236](#) shows such a topology, where R2 connects to the R1 source on one subnet, and to the incoming interface on R3 (ge-1/3/0.0 in the figure) on another subnet.

Figure 43: Accepting Multicast Traffic from a Remote Source



In this topology R2 is a pass-through device not running PIM, so R3 is the first hop router for multicast packets sent from R1. Because R1 and R3 are in different subnets, the default behavior of R3 is to disregard R1 as a remote source. You can have R3 accept multicast traffic from R1, however, by enabling **accept-remote-source** on the target interface.

To accept traffic from a remote source:

1. Identify the router and physical interface that you want to receive multicast traffic from the remote source.
2. Configure the interface to accept traffic from the remote source.

```
[edit protocols pim interface ge-1/3/0.0]
```



```
user@host# set accept-remote-source
```



NOTE: If the interface you identified is not the only path from the remote source, you need to ensure that it is the best path. For example you can configure a static route on the receiver side PE router to the source, or you can prepend the AS path on the other possible routes:

```
[edit policy-options policy-statement as-path-prepend term prepend]
user@host# set from route-filter 192.168.0.0/16 orlonger
user@host# set from route-filter 172.16.0.0/16 orlonger
user@host# set then as-path-prepend "1 1 1 1"
```

3. Commit the configuration changes.
4. Confirm that the interface you configured accepts traffic from the remote source.

```
user@host# show pim statistics
```

Related Documentation

- *Example: Allowing MBGP MVPN Remote Sources*
- *Understanding Prepending AS Numbers to BGP AS Paths*
- [show pim statistics on page 555](#)

Example: Configuring MSDP

Configure a router to act as a PIM sparse-mode rendezvous point and an MSDP peer:

```
[edit]
routing-options {
  interface-routes {
    rib-group ifrg;
  }
  rib-groups {
    ifrg {
      import-rib [inet.0 inet.2];
    }
    mcrg {
      export-rib inet.2;
      import-rib inet.2;
    }
  }
}
protocols {
  bgp {
    group lab {
      type internal;
      family any;
      neighbor 192.168.6.18 {
        local-address 192.168.6.17;
      }
    }
  }
}
```

```
pim {
  dense-groups {
    224.0.1.39/32;
    224.0.1.40/32;
  }
  rib-group mcr;
  rp {
    local {
      address 192.168.1.1;
    }
  }
  interface all {
    mode sparse-dense;
    version 1;
  }
}
msdp {
  rib-group mcr;
  group lab {
    peer 192.168.6.18 {
      local-address 192.168.6.17;
    }
  }
}
```

Example: Configuring MSDP with Active Source Limits and Mesh Groups

This example shows how to configure MSDP to filter source-active messages and limit the flooding of source-active messages.

- [Requirements on page 238](#)
- [Overview on page 238](#)
- [Configuration on page 242](#)
- [Verification on page 243](#)

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable PIM sparse mode. See “[PIM Overview](#)” on [page 125](#).
- Configure the router as a PIM sparse-mode RP. See “[Configuring Local PIM RPs](#)” on [page 181](#).

Overview

A router interested in MSDP messages, such as an RP, might have to process a large number of MSDP messages, especially source-active messages, arriving from other

routers. Because of the potential need for a router to examine, process, and create state tables for many MSDP packets, there is a possibility of an MSDP-based denial-of-service (DoS) attack on a router running MSDP. To minimize this possibility, you can configure the router to limit the number of source active messages the router accepts. Also, you can configure a threshold for applying random early detection (RED) to drop some but not all MSDP active source messages.

By default, the router accepts 25,000 source active messages before ignoring the rest. The limit can be from 1 through 1,000,000. The limit is applied to both the number of messages and the number of MSDP peers.

By default, the router accepts 24,000 source-active messages before applying the RED profile to prevent a possible DoS attack. This number can also range from 1 through 1,000,000. The next 1000 messages are screened by the RED profile and the accepted messages processed. If you configure no drop profiles (as this example does not), RED is still in effect and functions as the primary mechanism for managing congestion. In the default RED drop profile, when the packet queue fill-level is 0 percent, the drop probability is 0 percent. When the fill-level is 100 percent, the drop probability is 100 percent.



NOTE: The router ignores source-active messages with encapsulated TCP packets. Multicast does not use TCP; segments inside source-active messages are most likely the result of worm activity.

The number configured for the threshold must be less than the number configured for the maximum number of active MSDP sources.

You can configure an active source limit globally, for a group, or for a peer. If active source limits are configured at multiple levels of the hierarchy (as shown in this example), all are applied.

You can configure an active source limit for an address range as well as for a specific peer. A per-source active source limit uses an IP prefix and prefix length instead of a specific address. You can configure more than one per-source active source limit. The longest match determines the limit.

Per-source active source limits can be combined with active source limits at the peer, group, and global (instance) hierarchy level. Per-source limits are applied before any other type of active source limit. Limits are tested in the following order:

- Per-source
- Per-peer or group
- Per-instance

An active source message must “pass” all limits established before being accepted. For example, if a source is configured with an active source limit of 10,000 active multicast groups and the instance is configured with a limit of 5000 (and there are no other sources or limits configured), only 5000 active source messages are accepted from this source.

MSDP mesh groups are groups of peers configured in a full-mesh topology that limits the flooding of source-active messages to neighboring peers. Every mesh group member must have a peer connection with every other mesh group member. When a source-active message is received from a mesh group member, the source-active message is always accepted but is not flooded to other members of the same mesh group. However, the source-active message is flooded to non-mesh group peers or members of other mesh groups. By default, standard flooding rules apply if **mesh-group** is not specified.



CAUTION: When configuring MSDP mesh groups, you must configure all members the same way. If you do not configure a full mesh, excessive flooding of source-active messages can occur.

A common application for MSDP mesh groups is peer-reverse-path-forwarding (peer-RPF) check bypass. For example, if there are two MSDP peers inside an autonomous system (AS), and only one of them has an external MSDP session to another AS, the internal MSDP peer often rejects incoming source-active messages relayed by the peer with the external link. Rejection occurs because the external MSDP peer must be reachable by the internal MSDP peer through the next hop toward the source in another AS, and this next-hop condition is not certain. To prevent rejections, configure an MSDP mesh group on the internal MSDP peer so it always accepts source-active messages.



NOTE: An alternative way to bypass the peer-RPF check is to configure a default peer. In networks with only one MSDP peer, especially stub networks, the source-active message always needs to be accepted. An MSDP default peer is an MSDP peer from which all source-active messages are accepted without performing the peer-RPF check. You can establish a default peer at the peer or group level by including the **default-peer** statement.


Table 11 on page 240 explains how flooding is handled by peers in this example. .

Table 11: Source-Active Message Flooding Explanation

Source-Active Message Received From	Source-Active Message Flooded To	Source-Active Message Not Flooded To
Peer 21	Peer 11, Peer 12, Peer 13, Peer 31, Peer 32	Peer 22
Peer 11	Peer 21, Peer 22, Peer 31, Peer 32	Peer 12, Peer 13
Peer 31	Peer 21, Peer 22, Peer 11, Peer 12, Peer 13, Peer 32	—

Figure 44 on page 241 illustrates source-active message flooding between different mesh groups and peers within the same mesh group.

Figure 44: Source-Active Message Flooding



This example includes the following settings:

- **active-source-limit maximum 10000**—Applies a limit of 10,000 active sources to all other peers.
- **data-encapsulation disable**—On an RP router using MSDP, disables the default encapsulation of multicast data received in MSDP register messages inside MSDP source-active messages.

MSDP data encapsulation mainly concerns bursty sources of multicast traffic. Sources that send only one packet every few minutes have trouble with the timeout of state relationships between sources and their multicast groups (S,G). Routers lose data while they attempt to reestablish (S,G) state tables. As a result, multicast register messages contain data, and this data encapsulation in MSDP source-active messages can be turned on or off through configuration.

By default, MSDP data encapsulation is enabled. An RP running MSDP takes the data packets arriving in the source's register message and encapsulates the data inside an MSDP source-active message.

However, data encapsulation creates both a multicast forwarding cache entry in the **inet.1** table (this is also the forwarding table) and a routing table entry in the **inet.4** table. Without data encapsulation, MSDP creates only a routing table entry in the **inet.4** table. In some circumstances, such as the presence of Internet worms or other forms of DoS attack, the router's forwarding table might fill up with these entries. To prevent the forwarding table from filling up with MSDP entries, you can configure the router not to use MSDP data encapsulation. However, if you disable data encapsulation, the router ignores and discards the encapsulated data. Without data encapsulation, multicast applications with bursty sources having transmit intervals greater than about 3 minutes might not work well.

- **group MSDP-group local-address 10.1.2.3**—Specifies the address of the local router (this router).
- **group MSDP-group mode mesh-group**—Specifies that all peers belonging to the MSDP-group group are mesh group members.

- **group MSDP-group peer 10.10.10.10**—Prevents the sending of source-active messages to neighboring peer 10.10.10.10.
- **group MSDP-group peer 10.10.10.10 active-source-limit maximum 7500**—Applies a limit of 7500 active sources to MSDP peer 10.10.10.10 in group **MSDP-group**.
- **peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000**—Applies a threshold of 4000 active sources and a limit of 5000 active sources to MSDP peer 10.0.0.1.
- **source 10.1.0.0/16 active-source-limit maximum 500**—Applies a limit of 500 active sources to any source on the 10.1.0.0/16 network.

Configuration

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

```
set protocols msdp data-encapsulation disable
set protocols msdp active-source-limit maximum 10000
set protocols msdp peer 10.0.0.1 active-source-limit maximum 5000
set protocols msdp peer 10.0.0.1 active-source-limit threshold 4000
set protocols msdp source 10.1.0.0/16 active-source-limit maximum 500
set protocols msdp group MSDP-group mode mesh-group
set protocols msdp group MSDP-group local-address 10.1.2.3
set protocols msdp group MSDP-group peer 10.10.10.10 active-source-limit maximum
7500
```

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

To configure MSDP source active routes and mesh groups:

1. (Optional) Disable data encapsulation.

```
[edit protocols msdp]
user@host# set data-encapsulation disable
```

2. Configure the active source limits.

```
[edit protocols msdp]
user@host# set peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000
user@host# set group MSDP-group peer 10.10.10.10 active-source-limit maximum
7500
user@host# set active-source-limit maximum 10000
user@host# set source 10.1.0.0/16 active-source-limit maximum 500
```

3. Configure the mesh group.

```
[edit protocols msdp]
user@host# set group MSDP-group mode mesh-group
user@host# set group MSDP-group peer 10.10.10.10
user@host# set group MSDP-group local-address 10.1.2.3
```

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

```
[edit routing-instances]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols** command.

```
user@host# show protocols
msdp {
  data-encapsulation disable;
  active-source-limit {
    maximum 10000;
  }
  peer 10.0.0.1 {
    active-source-limit {
      maximum 5000;
      threshold 4000;
    }
  }
  source 10.1.0.0/16 {
    active-source-limit {
      maximum 500;
    }
  }
  group MSDP-group {
    mode mesh-group;
    local-address 10.1.2.3;
    peer 10.10.10.10 {
      active-source-limit {
        maximum 7500;
      }
    }
  }
}
```

Verification

To verify the configuration, run the following commands:

- [show msdp source-active](#)
- [show msdp statistics](#)

Related Documentation

- [Examples: Configuring MSDP](#)
- [Filtering MSDP SA Messages on page 234](#)
- [Configuring Local PIM RPs on page 181](#)

Example: Configuring PIM Anycast With or Without MSDP

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

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

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

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

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

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

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

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

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


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

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

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

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

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



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

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

```
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
```

```

        family inet {
            address 198.58.3.254/32 {
                primary;
            }
            address 198.58.3.253/32;
        }
    }
}

```

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

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

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

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

```

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

```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at

the `[edit protocols pim rp static]` hierarchy level. Include the **version** statement at the `[edit protocols pim rp static address]` hierarchy level to specify PIM version 2.

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

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

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

Configuring a PIM Anycast RP Router with MSDP

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

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

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

```
        interface fxp0.0 {  
            disable;  
        }  
    }  
}
```

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

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

PART 4

Configuration Statements and Operational Commands

- [Configuration Statements \(IGMP\) on page 251](#)
- [Configuration Statements \(IGMP Snooping\) on page 277](#)
- [Configuration Statements \(MLD Snooping\) on page 299](#)
- [Configuration Statements \(PIM\) on page 321](#)
- [Configuration Statements \(Source-Specific Multicast\) on page 395](#)
- [Configuration Statements \(MSDP\) on page 401](#)
- [Operational Commands \(IGMP\) on page 423](#)
- [Operational Commands \(IGMP Snooping\) on page 447](#)
- [Operational Commands \(PIM\) on page 459](#)
- [Operational Commands \(MSDP\) on page 565](#)

Configuration Statements (IGMP)

- [accounting \(Protocols IGMP\) on page 252](#)
- [accounting \(Protocols IGMP Interface\) on page 252](#)
- [asm-override-ssm on page 253](#)
- [disable \(Protocols IGMP\) on page 253](#)
- [exclude \(Protocols IGMP\) on page 254](#)
- [group \(Protocols IGMP\) on page 255](#)
- [group-count on page 256](#)
- [group-increment \(Protocols IGMP\) on page 256](#)
- [group-limit \(Protocols IGMP\) on page 257](#)
- [group-policy \(Protocols IGMP\) on page 257](#)
- [igmp on page 258](#)
- [immediate-leave \(Protocols IGMP\) on page 260](#)
- [interface \(Protocols IGMP\) on page 261](#)
- [maximum-transmit-rate \(Protocols IGMP\) on page 262](#)
- [oif-map \(IGMP Interface\) on page 262](#)
- [passive \(IGMP\) on page 263](#)
- [promiscuous-mode \(Protocols IGMP\) on page 264](#)
- [query-interval \(Protocols IGMP\) on page 265](#)
- [query-last-member-interval \(Protocols IGMP\) on page 266](#)
- [query-response-interval \(Protocols IGMP\) on page 267](#)
- [robust-count \(Protocols IGMP\) on page 268](#)
- [source \(Protocols IGMP\) on page 269](#)
- [source-count \(Protocols IGMP\) on page 270](#)
- [source-increment \(Protocols IGMP\) on page 271](#)
- [static \(Protocols IGMP\) on page 272](#)
- [traceoptions \(Protocols IGMP\) on page 273](#)
- [version \(Protocols IGMP\) on page 275](#)

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.
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">• Recording IGMP Join and Leave Events on page 59

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.
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">• Recording IGMP Join and Leave Events on page 59

asm-override-ssm

Syntax	asm-override-ssm;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Enable the routing device to accept any-source multicast join messages (*G) for group addresses that are within the default or configured range of source-specific multicast groups.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 172


disable (Protocols IGMP)

Syntax	disable;
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 before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Disable IGMP on the system.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Disabling IGMP on page 63

exclude (Protocols IGMP)

Syntax	exclude;
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 in Junos OS Release 9.3.
Description	Configure the static group to operate in exclude mode. In exclude mode all sources except the address configured are accepted for the group. If this statement is not included, the group operates in include mode.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling IGMP Static Group Membership on page 52

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.
Description	Specify the IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.
<div>  NOTE: You must specify a unique address for each group. </div>	
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"> • Enabling IGMP Static Group Membership on page 52

group-count

Syntax	<code>group-count <i>number</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 in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the number of static groups to be created.
Options	<i>number</i> —Number of static groups. Default: Range: 1 through 512
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling IGMP Static Group Membership on page 52

group-increment (Protocols IGMP)

Syntax	<code>group-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 in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the number of times the address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.
Options	<i>increment</i> —Number of times the address should be incremented. Default: 0.0.0.1 Range: 0.0.0.1 through 255.255.255.255
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling IGMP Static Group Membership on page 52

group-limit (Protocols IGMP)

Syntax	<code>group-limit <i>limit</i>;</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 10.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.
Default	By default, there is no limit to the number of multicast groups that can join the interface.
Options	<i>limit</i> —group limit value for the interface. Range: 1 through 32767
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 60

group-policy (Protocols IGMP)

Syntax	<code>group-policy [<i>policy-names</i>];</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 9.1. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the routing device receives an IGMP report, the routing device compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 48

igmp

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

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Enable IGMP on the router. IGMP must be enabled for the router to receive multicast packets.


The remaining statements are explained separately.

Default IGMP is disabled on the router. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

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

Related Documentation • [Enabling IGMP on page 43](#)

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.
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>The immediate leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</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 router only knows about the one interested host and does not know about the others.</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"> Specifying Immediate-Leave Host Removal for IGMP on page 47

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	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable IGMP on an interface and configure interface-specific properties.
Options	<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 on page 43


maximum-transmit-rate (Protocols IGMP)

Syntax	maximum-transmit-rate <i>packets-per-second</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols igmp], [edit protocols igmp]
Release Information	Statement introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Limit the transmission rate of IGMP packets
Options	packets-per-second —Maximum number of IGMP packets transmitted in one second by the router. Range: 1 through 10000 Default: 500 packets
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Limiting the Maximum IGMP Message Rate on page 52

oif-map (IGMP Interface)

Syntax	oif-map <i>map-name</i> ;
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 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Associates an outgoing interface (OIF) map to the IGMP interface. The OIF map is a routing policy statement that can contain multiple terms.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Multicast with Subscriber VLANs

passive (IGMP)

Syntax	<code>passive <allow-receive> <send-general-query> <send-group-query>;</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 9.6. allow-receive , send-general-query , and send-group-query options were added in Junos OS Release 10.0. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify that IGMP run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as IGMP reports, queries, and leaves.
<div>  <p>NOTE: You can selectively activate up to two out of the three available options for the passive statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the passive statement.</p> </div>	
Options	<p>allow-receive—Enables IGMP to receive control traffic on the interface.</p> <p>send-general-query—Enables IGMP to send general queries on the interface.</p> <p>send-group-query—Enables IGMP to send group-specific and group-source-specific queries on the interface.</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 Multicast with Subscriber VLANs</i> • Enabling IGMP on page 43

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 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 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>• Accepting IGMP Messages from Remote Subnetworks on page 49

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"> • Modifying the IGMP Host-Query Message Interval on page 45 • query-last-member-interval (Protocols IGMP) on page 266 • query-response-interval (Protocols IGMP) on page 267

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	seconds —Time interval, in fractions of a second or seconds. Range: 0.1 through 0.9, then in 1-second intervals 1 through 999999 Default: 1 second
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Modifying the IGMP Last-Member Query Interval on page 46• query-interval (Protocols IGMP) on page 265• query-response-interval (Protocols IGMP) on page 267

query-response-interval (Protocols IGMP)

Syntax	<code>query-response-interval <i>seconds</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	Specify how long the querier routing device waits to receive a response to a host-query message from a host.
Options	<i>seconds</i> —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"> • Modifying the IGMP Query Response Interval on page 50 • query-interval (Protocols IGMP) on page 265 • query-last-member-interval (Protocols IGMP) on page 266

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 on page 51

source (Protocols IGMP)

Syntax	<pre>source <i>ip-address</i> { <i>source-count</i> <i>number</i>; <i>source-increment</i> <i>increment</i>; }</pre>
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	<p><i>ip-address</i>—IPv4 unicast address.</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"> • Enabling IGMP Static Group Membership on page 52


source-count (Protocols IGMP)

Syntax	<code>source-count <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols igmp interface <i>interface-name</i> static group multicast-group-address source], [edit protocols igmp interface <i>interface-name</i> static group multicast-group-address source]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the number of multicast source addresses that should be accepted for each static group created.
Options	<i>number</i> —Number of source addresses. Default: 1 Range: 1 through 1024
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling IGMP Static Group Membership on page 52

source-increment (Protocols IGMP)

Syntax	source-increment <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols igmp interface <i>interface-name</i> static group multicast-group-address source], [edit protocols igmp interface <i>interface-name</i> static group multicast-group-address source]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the number of times the multicast source address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.
Options	increment —Number of times the source address should be incremented. Default: 0.0.0.1 Range: 0.0.0.1 through 255.255.255.255
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling IGMP Static Group Membership on page 52

static (Protocols IGMP)

Syntax	<pre>static { group multicast-group-address { exclude; group-count number; group-increment increment; source ip-address { source-count number; source-increment increment; } } }</pre>
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	<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	<p>Test multicast forwarding on an interface without a receiver host.</p> <p>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.</p>
<div>  <p>NOTE: To prevent joining too many groups accidentally, the static statement is not supported with the interface all statement.</p> </div>	
The remaining statements are explained separately.	
Required Privilege Level	<p>routing and trace—To view this statement in the configuration.</p> <p>routing-control and trace-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling IGMP Static Group Membership on page 52

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	<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	<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 Level	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Tracing IGMP Protocol Traffic on page 61

version (Protocols IGMP)

Syntax	<code>version <i>version</i>;</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 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	<p>version—IGMP version number.</p> <p>Range: 1, 2, or 3</p> <p>Default: IGMP version 2</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"> • Changing the IGMP Version on page 44

CHAPTER 18

Configuration Statements (IGMP Snooping)

- [all](#) on page 278
- [data-forwarding](#) on page 278
- [disable \(IGMP Snooping\)](#) on page 279
- [group \(IGMP Snooping\)](#) on page 279
- [group-limit \(IGMP and MLD Snooping\)](#) on page 280
- [host-only-interface](#) on page 281
- [igmp-querier](#) on page 281
- [igmp-snooping](#) on page 282
- [immediate-leave \(Bridge Domains\)](#) on page 283
- [interface \(Bridge Domains\)](#) on page 284
- [interface \(IGMP Snooping\)](#) on page 285
- [l2-querier](#) on page 285
- [multicast-router-interface \(IGMP Snooping\)](#) on page 286
- [query-interval \(Bridge Domains\)](#) on page 287
- [query-last-member-interval \(Bridge Domains\)](#) on page 288
- [query-response-interval \(Bridge Domains\)](#) on page 289
- [receiver](#) on page 290
- [robust-count \(IGMP Snooping\)](#) on page 290
- [source-address](#) on page 291
- [src-address \(IGMP Querier\)](#) on page 292
- [source-vlans](#) on page 292
- [static \(IGMP Snooping\)](#) on page 293
- [traceoptions \(IGMP Snooping\)](#) on page 294
- [version \(IGMP Snooping\)](#) on page 296
- [vlan \(IGMP Snooping\)](#) on page 297

all

Syntax	all;
Hierarchy Level	[edit protocols igmp-snooping vlan]
Release Information	Statement introduced in Junos OS Release 15.1 for the QFX series.
Description	Apply IGMP snooping to all configured VLANs.
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 on page 68

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">• <i>Example: Configuring Multicast VLAN Registration</i>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i>

disable (IGMP Snooping)

Syntax	<code>disable;</code>
Hierarchy Level	<code>[edit protocols igmp-snooping vlan <i>vlan-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Disable IGMP snooping on all interfaces in a VLAN.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IGMP Snooping on page 70 • Configuring IGMP Snooping on page 68

group (IGMP Snooping)

Syntax	<code>group <i>ip-address</i>;</code>
Hierarchy Level	<code>[edit protocols igmp-snooping vlan <i>vlan-name</i> interface <i>interface-name</i> static]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a static multicast group using a valid IP multicast address.
Default	None.
Options	<i>ip-address</i> —IP address of the multicast group receiving data on an interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show igmp-snooping vlans on page 457 • Example: Configuring IGMP Snooping on page 70 • Configuring IGMP Snooping on page 68

group-limit (IGMP and MLD Snooping)

Syntax	<code>group-limit <i>limit</i>;</code>
Hierarchy Level	<code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>],</code> <code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>],</code> <code>[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>]</code>
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.
Default	By default, there is no limit to the number of multicast groups joining an interface.
Options	<i>limit</i> —a 32-bit number for the limit on the interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IGMP Snooping</i>

host-only-interface

Syntax	host-only-interface;
Hierarchy Level	[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface interface-name], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface interface-name], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface interface-name], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping interface interface-name]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Configure an interface as a host-facing interface. IGMP queries received on these interfaces are dropped.
Default	The interface can either be a host-side or multicast-router 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"> • Example: Configuring IGMP Snooping • multicast-router-interface

igmp-querier

Syntax	igmp-querier source-address <i>source address</i> ;
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 14.1X53-D15 for QFabric Systems.
Description	Configure a QFabric Node device to be an IGMP querier. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is lower (a smaller number) than the IP addresses for those routers on the network. This ensures that Node is always the IGMP querier on the network.
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 IGMP Snooping on page 70 • Configuring IGMP Snooping on page 68 • show igmp-snooping vlans on page 457 • show configuration protocols igmp on page 433

igmp-snooping

Syntax	<pre> igmp-snooping { vlan <i>vlan-id</i> { all immediate-leave; interface <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>; } } } } l2-querier { source-address <i>ip-address</i>; } proxy { source-address <i>ip-address</i>; } 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>> <no-stamp> <replace> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier>; } } </pre>
Hierarchy Level	[edit protocols]
Release Information	Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	Enable IGMP snooping on the router or switch.
Default	IGMP snooping is disabled on the router or switch.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding IGMP Snooping</i> • <i>IGMP Snooping in MC-LAG Active-Active Mode</i>

immediate-leave (Bridge Domains)

Syntax	<code>immediate-leave;</code>
Hierarchy Level	<pre>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>], [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>]</pre>
Release Information	Statement introduced in Junos OS Release 8.5.
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>The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</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>



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

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*

interface (Bridge Domains)

Syntax

```
interface interface-name {  
    group-limit limit;  
    host-only-interface;  
    multicast-router-interface;  
    static {  
        group ip-address {  
            source ip-address;  
        }  
    }  
}
```

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols
igmp-snooping],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols
vlan *vlan-id* igmp-snooping]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable IGMP snooping 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

- *Example: Configuring IGMP Snooping*

interface (IGMP Snooping)

Syntax	<pre>interface <i>interface-name</i> { multicast-router-interface; static { group <i>ip-address</i>; } }</pre>
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Enable IGMP snooping on an interface and configure interface-specific properties.</p> <p>The remaining statements are explained separately.</p>
Options	<i>interface-name</i> —Name of the 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"> • Example: Configuring IGMP Snooping on page 70 • Configuring IGMP Snooping on page 68 • show igmp-snooping vlans on page 457

l2-querier

Syntax	<pre>l2-querier { source-address <i>ip-address</i>; }</pre>
Hierarchy Level	[edit protocols igmp-snooping vlan],
Release Information	Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	Configure the switch to be an IGMP querier. Use the source-address statement to configure the source address to use for IGMP snooping queries.
Options	<p>seconds—Time interval.</p> <p>Range: 1 through 1024</p> <p>Default: 125 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	

multicast-router-interface (IGMP Snooping)

Syntax	multicast-router-interface;
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure an interface to forward IGMP messages to multicast routers.
Default	Disabled. If this statement is disabled, the interface drops IGMP messages it receives.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show igmp-snooping vlans on page 457• Example: Configuring IGMP Snooping on page 70• Configuring IGMP Snooping on page 68

query-interval (Bridge Domains)

Syntax	<code>query-interval seconds;</code>
Hierarchy Level	<pre>[edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping] , [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <i>interface</i> <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <i>interface</i> <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <i>interface</i> <i>interface-name</i>],[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols mld-snooping] [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <i>interface</i> <i>interface-name</i>] [edit routing-instances <i>routing-instance-name</i> protocols mld-snooping] [edit protocols igmp-snooping vlan]</pre>
Release Information	<p>Statement introduced before Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.</p>
Description	Configure the interval for host-query message timeouts.
Options	<p><i>seconds</i>—Time interval. This value must be greater than the interval set for <i>query-response-interval</i>.</p> <p>Range: 1 through 1024</p> <p>Default: 125 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>Example: Configuring IGMP Snooping</i> • query-last-member-interval (Bridge Domains) on page 288 • query-response-interval (Bridge Domains) on page 289 • <i>mld-snooping</i> • <i>igmp-snooping</i>

query-last-member-interval (Bridge Domains)

Syntax	<code>query-last-member-interval seconds;</code>
Hierarchy Level	<pre>[edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping] , [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols mld-snooping] [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> [edit routing-instances <i>routing-instance-name</i> protocols mld-snooping]interface <i>interface-name</i>] [edit protocols igmp-snooping vlan],</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.</p>
Description	Configure the interval for group-specific query timeouts.
Options	<p>seconds—Time interval, in fractions of a second or seconds.</p> <p>Range: 0.1 through 0.9, then in 1-second intervals 1 through 1024</p> <p>Default: 1 second</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 IGMP Snooping</i> • query-interval on page 287 • query-response-interval on page 289 • <i>mld-snooping</i> • <i>igmp-snooping</i>

query-response-interval (Bridge Domains)

Syntax	<code>query-response-interval seconds;</code>
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-snoopingvlan <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping] ,</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 igmp-snoopingvlan <i>vlan-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols mld-snooping]interface <i>interface-name</i>]</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mld-snooping]</p> <p>[edit protocols igmp-snooping vlan],</p>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.</p>
Description	Specify how long to wait to receive a response to a specific query message from a host.
Options	<p><i>seconds</i>—Time interval. This interval should be less than the host-query interval.</p> <p>Range: 1 through 1024</p> <p>Default: 10 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>Example: Configuring IGMP Snooping</i> • query-interval (Bridge Domains) on page 287 • query-last-member-interval (Bridge Domains) on page 288 • <i>mld-snooping</i> • <i>igmp-snooping</i>

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	Configure a VLAN as a multicast receiver VLAN of the multicast VLAN (MVLAN). The remaining statements are explained separately.
Default	Disabled
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Multicast VLAN Registration</i>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i>

robust-count (IGMP Snooping)

Syntax	<pre>robust-count <i>number</i>;</pre>
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the number of intervals the switch waits before removing a multicast group from the multicast forwarding table. Configure the length of each interval using the query-interval statement.
Default	2 intervals
Options	<i>number</i> —Number of intervals the switch waits before timing out a multicast group. Range: 2 through 10
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IGMP Snooping on page 70• Configuring IGMP Snooping on page 68• show igmp-snooping vlans on page 457

source-address

Syntax	<code>source-address <i>ip-address</i>;</code>
Hierarchy Level	<p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping proxy],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> proxy],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping proxy],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> proxy]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p>
Description	<p>Specify the IP address to use as the source for IGMP snooping reports in proxy mode. Reports are sent with address 0.0.0.0 as the source address unless there is a source address configured. You can also use this statement to configure the source address to use for IGMP snooping queries.</p>
Options	<i>ip-address</i> —IP address to use as the source for proxy-mode IGMP snooping reports.
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>

src-address (IGMP Querier)

Syntax	<code>src-address source address;</code>
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i> igmp-querier] [edit protocols igmp-snooping vlan <i>vlan-name</i> l2-querier]
Release Information	Statement introduced in Junos OS Release 12.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D15 for QFabric Systems.
Description	Configure the address that the switch uses as the source address in the IGMP queries that it sends. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is smaller (a lower number) than the IP addresses for those routers on the network. This ensures that switch is always the IGMP querier on the network.
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 IGMP Snooping on page 70• Configuring IGMP Snooping on page 68• show igmp-snooping vlans on page 457• show configuration protocols igmp on page 433

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• Configuring Multicast VLAN Registration (CLI Procedure)

static (IGMP Snooping)

Syntax	<pre>static { group ip-address; }</pre>
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Statically define multicast groups on an interface.</p> <p>The remaining statement is explained separately.</p>
Default	No multicast groups are statically defined.
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 page 70• Configuring IGMP Snooping on page 68• show igmp-snooping vlans on page 457

traceoptions (IGMP Snooping)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <size <i>size</i>> <replace> <world-readable no-world-readable>; flag <i>flag</i> (detail disable receive send); }</pre>
Hierarchy Level	For platforms without ELS: [edit protocols igmp-snooping] For platforms with ELS: [edit protocols igmp-snooping vlan]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define tracing operations for IGMP snooping.
Default	The traceoptions feature is disabled by default.
Options	<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. 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 (xk to specify KB, xm to specify MB, or xg to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 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 sockets.• nexthop— Trace next-hop related events.• normal—Trace normal IGMP snooping protocol events.• packets—Trace all IGMP packets.• policy—Trace policy processing.• query—Trace IGMP membership query messages.• report—Trace membership report messages.

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

no-stamp—(Optional) Do not time stamp trace file.

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

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

Range: 10 KB through 1 gigabytes

Default: 128 KB

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

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IGMP Snooping on page 70 • Configuring IGMP Snooping on page 68

version (IGMP Snooping)

Syntax	<code>version <i>number</i>;</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)• Configuring IGMP Snooping on page 68

vlan (IGMP Snooping)

Syntax	<pre> vlan <i>vlan-name</i> { <i>immediate-leave</i>; interface <i>interface-name</i> { <i>group-limit limit</i>; <i>host-only-interface</i>; multicast-router-interface; static { group <i>mcast-group-address</i> { source <i>ip-address</i>; } } } <i>qualified-vlan</i> ; proxy { <i>source-address ip-address</i>; } <i>query-interval seconds</i>; <i>query-last-member-interval seconds</i>; <i>query-response-interval seconds</i>; <i>robust-count number</i>; } </pre>
Hierarchy Level	[edit protocols igmp-snooping],
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 13.2 for the QFX series.
Description	Configure IGMP snooping parameters for a particular VLAN.
Default	By default, IGMP snooping options apply to all VLANs.
Options	<p><i>vlan-name</i>—Apply the parameters to this VLAN.</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 VLAN-Specific IGMP Snooping Parameters on page 69 igmp-snooping on page 282

CHAPTER 19

Configuration Statements (MLD Snooping)

- all on page 300
- data-forwarding on page 300
- disable (IGMP Snooping) on page 301
- group (IGMP Snooping) on page 301
- group-limit (IGMP and MLD Snooping) on page 302
- host-only-interface on page 303
- igmp-querier on page 303
- igmp-snooping on page 304
- immediate-leave (Bridge Domains) on page 305
- interface (Bridge Domains) on page 306
- interface (IGMP Snooping) on page 307
- l2-querier on page 307
- mld-snooping on page 308
- multicast-router-interface (IGMP Snooping) on page 309
- query-interval (Bridge Domains) on page 310
- query-last-member-interval (Bridge Domains) on page 311
- query-response-interval (Bridge Domains) on page 312
- receiver on page 313
- robust-count (IGMP Snooping) on page 313
- source-address on page 314
- src-address (IGMP Querier) on page 315
- source-vlans on page 315
- static (IGMP Snooping) on page 316
- traceoptions (IGMP Snooping) on page 317
- version (IGMP Snooping) on page 319
- vlan (IGMP Snooping) on page 320

all

Syntax	all;
Hierarchy Level	[edit protocols igmp-snooping vlan]
Release Information	Statement introduced in Junos OS Release 15.1 for the QFX series.
Description	Apply IGMP snooping to all configured VLANs.
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 on page 68

data-forwarding

Syntax	<pre>data-forwarding { receiver { source-vlans vlan-list; install; } source { groups group-prefix; } }</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"> • <i>Example: Configuring Multicast VLAN Registration</i> • <i>Configuring Multicast VLAN Registration (CLI Procedure)</i>

disable (IGMP Snooping)

Syntax	<code>disable;</code>
Hierarchy Level	<code>[edit protocols igmp-snooping vlan <i>vlan-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Disable IGMP snooping on all interfaces in a VLAN.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IGMP Snooping on page 70 • Configuring IGMP Snooping on page 68

group (IGMP Snooping)

Syntax	<code>group <i>ip-address</i>;</code>
Hierarchy Level	<code>[edit protocols igmp-snooping vlan <i>vlan-name</i> interface <i>interface-name</i> static]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a static multicast group using a valid IP multicast address.
Default	None.
Options	<i>ip-address</i> —IP address of the multicast group receiving data on an interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • show igmp-snooping vlans on page 457 • Example: Configuring IGMP Snooping on page 70 • Configuring IGMP Snooping on page 68

group-limit (IGMP and MLD Snooping)

Syntax	<code>group-limit <i>limit</i>;</code>
Hierarchy Level	<code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>],</code> <code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>],</code> <code>[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>]</code>
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.
Default	By default, there is no limit to the number of multicast groups joining an interface.
Options	<i>limit</i> —a 32-bit number for the limit on the interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IGMP Snooping</i>

host-only-interface

Syntax	host-only-interface;
Hierarchy Level	[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface interface-name], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface interface-name], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface interface-name], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping interface interface-name]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Configure an interface as a host-facing interface. IGMP queries received on these interfaces are dropped.
Default	The interface can either be a host-side or multicast-router 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"> • Example: Configuring IGMP Snooping • multicast-router-interface

igmp-querier

Syntax	igmp-querier source-address <i>source address</i> ;
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 14.1X53-D15 for QFabric Systems.
Description	Configure a QFabric Node device to be an IGMP querier. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is lower (a smaller number) than the IP addresses for those routers on the network. This ensures that Node is always the IGMP querier on the network.
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 IGMP Snooping on page 70 • Configuring IGMP Snooping on page 68 • show igmp-snooping vlans on page 457 • show configuration protocols igmp on page 433

igmp-snooping

Syntax	<pre> igmp-snooping { vlan <i>vlan-id</i> { all immediate-leave; interface <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>; } } } } l2-querier { source-address <i>ip-address</i>; } proxy { source-address <i>ip-address</i>; } 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>> <no-stamp> <replace> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier>; } } </pre>
Hierarchy Level	[edit protocols]
Release Information	Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	Enable IGMP snooping on the router or switch.
Default	IGMP snooping is disabled on the router or switch.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding IGMP Snooping</i> • <i>IGMP Snooping in MC-LAG Active-Active Mode</i>

immediate-leave (Bridge Domains)

Syntax	<code>immediate-leave;</code>
Hierarchy Level	<pre>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>], [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>]</pre>
Release Information	Statement introduced in Junos OS Release 8.5.
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>The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</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>



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

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*

interface (Bridge Domains)

Syntax

```
interface interface-name {  
    group-limit limit;  
    host-only-interface;  
    multicast-router-interface;  
    static {  
        group ip-address {  
            source ip-address;  
        }  
    }  
}
```

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable IGMP snooping 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

- *Example: Configuring IGMP Snooping*

interface (IGMP Snooping)

Syntax	<pre>interface <i>interface-name</i> { multicast-router-interface; static { group <i>ip-address</i>; } }</pre>
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Enable IGMP snooping on an interface and configure interface-specific properties.</p> <p>The remaining statements are explained separately.</p>
Options	<i>interface-name</i> —Name of the 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"> • Example: Configuring IGMP Snooping on page 70 • Configuring IGMP Snooping on page 68 • show igmp-snooping vlans on page 457

l2-querier

Syntax	<pre>l2-querier { source-address <i>ip-address</i>; }</pre>
Hierarchy Level	[edit protocols igmp-snooping vlan],
Release Information	Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	Configure the switch to be an IGMP querier. Use the source-address statement to configure the source address to use for IGMP snooping queries.
Options	<p>seconds—Time interval.</p> <p>Range: 1 through 1024</p> <p>Default: 125 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	

mld-snooping

Syntax	<pre> mld-snooping { vlan (<i>vlan-name</i>) { 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 <i>vlan-id</i>; 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> <<i>flag-modifier</i>>; } } </pre>
Hierarchy Level	[edit protocols] [edit routing-instances <i>instance-name</i> protocols]
Release Information	Statement introduced in Junos OS Release 13.3 for EX Series switches.
Description	<p>Enable and configure MLD snooping.</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"> Example: Configuring MLD Snooping on page 115 Configuring MLD Snooping on a VLAN (CLI Procedure) on page 109

multicast-router-interface (IGMP Snooping)

Syntax	<code>multicast-router-interface;</code>
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure an interface to forward IGMP messages to multicast routers.
Default	Disabled. If this statement is disabled, the interface drops IGMP messages it receives.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show igmp-snooping vlans on page 457• Example: Configuring IGMP Snooping on page 70• Configuring IGMP Snooping on page 68

query-interval (Bridge Domains)

Syntax	<code>query-interval seconds;</code>
Hierarchy Level	<pre>[edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping] , [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>],[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols mld-snooping] [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>] [edit routing-instances <i>routing-instance-name</i> protocols mld-snooping] [edit protocols igmp-snooping vlan]</pre>
Release Information	<p>Statement introduced before Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.</p>
Description	Configure the interval for host-query message timeouts.
Options	<p><i>seconds</i>—Time interval. This value must be greater than the interval set for <code>query-response-interval</code>.</p> <p>Range: 1 through 1024</p> <p>Default: 125 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>Example: Configuring IGMP Snooping</i> • query-last-member-interval (Bridge Domains) on page 288 • query-response-interval (Bridge Domains) on page 289 • <i>mld-snooping</i> • <i>igmp-snooping</i>

query-last-member-interval (Bridge Domains)

Syntax	<code>query-last-member-interval seconds;</code>
Hierarchy Level	<pre>[edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping] , [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols mld-snooping] [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> [edit routing-instances <i>routing-instance-name</i> protocols mld-snooping]interface <i>interface-name</i>] [edit protocols igmp-snooping vlan],</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.</p>
Description	Configure the interval for group-specific query timeouts.
Options	<p>seconds—Time interval, in fractions of a second or seconds.</p> <p>Range: 0.1 through 0.9, then in 1-second intervals 1 through 1024</p> <p>Default: 1 second</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 IGMP Snooping</i> • query-interval on page 287 • query-response-interval on page 289 • <i>mld-snooping</i> • <i>igmp-snooping</i>

query-response-interval (Bridge Domains)

Syntax	<code>query-response-interval seconds;</code>
Hierarchy Level	<pre>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snoopingvlan <i>vlan-id</i> interface <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols mld-snooping] , [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snoopingvlan <i>vlan-id</i> [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols mld-snooping]interface <i>interface-name</i>] [edit routing-instances <i>routing-instance-name</i> protocols mld-snooping] [edit protocols igmp-snooping vlan],</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.2 for MX series Routers with MPC.</p>
Description	Specify how long to wait to receive a response to a specific query message from a host.
Options	<p><i>seconds</i>—Time interval. This interval should be less than the host-query interval.</p> <p>Range: 1 through 1024</p> <p>Default: 10 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>Example: Configuring IGMP Snooping</i> • query-interval (Bridge Domains) on page 287 • query-last-member-interval (Bridge Domains) on page 288 • <i>mld-snooping</i> • <i>igmp-snooping</i>

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	Configure a VLAN as a multicast receiver VLAN of the multicast VLAN (MVLAN). The remaining statements are explained separately.
Default	Disabled
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Multicast VLAN Registration</i> • <i>Configuring Multicast VLAN Registration (CLI Procedure)</i>

robust-count (IGMP Snooping)

Syntax	robust-count <i>number</i> ;
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the number of intervals the switch waits before removing a multicast group from the multicast forwarding table. Configure the length of each interval using the query-interval statement.
Default	2 intervals
Options	<i>number</i> —Number of intervals the switch waits before timing out a multicast group. Range: 2 through 10
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IGMP Snooping on page 70 • Configuring IGMP Snooping on page 68 • show igmp-snooping vlans on page 457

source-address

Syntax	<code>source-address <i>ip-address</i>;</code>
Hierarchy Level	[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping proxy], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> proxy], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping proxy], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> proxy]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 13.2 for the QFX series.
Description	Specify the IP address to use as the source for IGMP snooping reports in proxy mode. Reports are sent with address 0.0.0.0 as the source address unless there is a source address configured. You can also use this statement to configure the source address to use for IGMP snooping queries.
Options	<i>ip-address</i> —IP address to use as the source for proxy-mode IGMP snooping reports.
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 IGMP Snooping</i>

src-address (IGMP Querier)

Syntax	<code>src-address source address;</code>
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i> igmp-querier] [edit protocols igmp-snooping vlan <i>vlan-name</i> l2-querier]
Release Information	Statement introduced in Junos OS Release 12.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D15 for QFabric Systems.
Description	Configure the address that the switch uses as the source address in the IGMP queries that it sends. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is smaller (a lower number) than the IP addresses for those routers on the network. This ensures that switch is always the IGMP querier on the network.
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 IGMP Snooping on page 70 • Configuring IGMP Snooping on page 68 • show igmp-snooping vlans on page 457 • show configuration protocols igmp on page 433

source-vlans

Syntax	<code>source-vlans vlan-list;</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 • Configuring Multicast VLAN Registration (CLI Procedure)

static (IGMP Snooping)

Syntax	<pre>static { group ip-address; }</pre>
Hierarchy Level	[edit protocols igmp-snooping vlan <i>vlan-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Statically define multicast groups on an interface.</p> <p>The remaining statement is explained separately.</p>
Default	No multicast groups are statically defined.
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 page 70• Configuring IGMP Snooping on page 68• show igmp-snooping vlans on page 457

traceoptions (IGMP Snooping)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <size <i>size</i>> <replace> <world-readable no-world-readable>; flag <i>flag</i> (detail disable receive send); } </pre>
Hierarchy Level	<p>For platforms without ELS:</p> <pre>[edit protocols igmp-snooping]</pre> <p>For platforms with ELS:</p> <pre>[edit protocols igmp-snooping vlan]</pre>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define tracing operations for IGMP snooping.
Default	The traceoptions feature is disabled by default.
Options	<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. 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 (xk to specify KB, xm to specify MB, or xg to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 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 sockets. • nexthop— Trace next-hop related events. • normal—Trace normal IGMP snooping protocol events. • packets—Trace all IGMP packets. • policy—Trace policy processing. • query—Trace IGMP membership query messages. • report—Trace membership report messages.

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

no-stamp—(Optional) Do not time stamp trace file.

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

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

Range: 10 KB through 1 gigabytes

Default: 128 KB

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

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	• Example: Configuring IGMP Snooping on page 70
	• Configuring IGMP Snooping on page 68

version (IGMP Snooping)

Syntax	<code>version <i>number</i>;</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) • Configuring IGMP Snooping on page 68

vlan (IGMP Snooping)

Syntax	<pre> vlan <i>vlan-name</i> { <i>immediate-leave</i>; interface <i>interface-name</i> { <i>group-limit limit</i>; <i>host-only-interface</i>; multicast-router-interface; static { group <i>multicast-group-address</i> { source <i>ip-address</i>; } } } <i>qualified-vlan</i> ; proxy { <i>source-address ip-address</i>; } <i>query-interval seconds</i>; <i>query-last-member-interval seconds</i>; <i>query-response-interval seconds</i>; <i>robust-count number</i>; } </pre>
Hierarchy Level	[edit protocols igmp-snooping],
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 13.2 for the QFX series.
Description	Configure IGMP snooping parameters for a particular VLAN.
Default	By default, IGMP snooping options apply to all VLANs.
Options	<p><i>vlan-name</i>—Apply the parameters to this VLAN.</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 VLAN-Specific IGMP Snooping Parameters on page 69 igmp-snooping on page 282

CHAPTER 20

Configuration Statements (PIM)

- [address \(Anycast RPs\) on page 323](#)
- [address \(Local RPs\) on page 324](#)
- [address \(Static RPs\) on page 325](#)
- [algorithm on page 326](#)
- [anycast-pim on page 327](#)
- [assert-timeout on page 328](#)
- [authentication on page 329](#)
- [auto-rp on page 330](#)
- [bfd-liveness-detection on page 331](#)
- [bootstrap on page 332](#)
- [bootstrap-export on page 333](#)
- [bootstrap-import on page 334](#)
- [bootstrap-priority on page 335](#)
- [dense-groups on page 336](#)
- [detection-time \(BFD for PIM\) on page 337](#)
- [disable \(PIM\) on page 338](#)
- [dr-election-on-p2p on page 339](#)
- [dr-register-policy on page 339](#)
- [embedded-rp on page 340](#)
- [export \(Bootstrap\) on page 341](#)
- [export \(Protocols PIM\) on page 342](#)
- [family \(Bootstrap\) on page 343](#)
- [family \(Protocols PIM\) on page 344](#)
- [family \(Local RP\) on page 345](#)
- [group \(RPF Selection\) on page 346](#)
- [group-ranges on page 347](#)
- [hello-interval on page 348](#)
- [hold-time \(Protocols PIM\) on page 349](#)

- [import \(Protocols PIM Bootstrap\) on page 350](#)
- [import \(Protocols PIM\) on page 351](#)
- [infinity on page 352](#)
- [interface on page 353](#)
- [join-load-balance on page 354](#)
- [join-prune-timeout on page 355](#)
- [key-chain on page 355](#)
- [local on page 356](#)
- [local-address \(Protocols PIM\) on page 357](#)
- [loose-check on page 358](#)
- [mapping-agent-election on page 359](#)
- [maximum-rps on page 360](#)
- [minimum-interval \(PIM BFD Liveness Detection\) on page 361](#)
- [minimum-interval \(PIM BFD Transmit Interval\) on page 362](#)
- [minimum-receive-interval on page 363](#)
- [mode \(Protocols PIM\) on page 363](#)
- [multiplier on page 364](#)
- [neighbor-policy on page 364](#)
- [next-hop \(PIM RPF Selection\) on page 365](#)
- [no-adaptation \(PIM BFD Liveness Detection\) on page 365](#)
- [override-interval on page 366](#)
- [pim on page 367](#)
- [prefix-list \(PIM RPF Selection\) on page 370](#)
- [priority \(Bootstrap\) on page 371](#)
- [priority \(PIM Interfaces\) on page 372](#)
- [priority \(PIM RPs\) on page 373](#)
- [propagation-delay on page 374](#)
- [register-probe-time on page 375](#)
- [reset-tracking-bit on page 376](#)
- [rib-group \(Protocols PIM\) on page 377](#)
- [rp on page 378](#)
- [rp-register-policy on page 380](#)
- [rp-set on page 381](#)
- [rpf-selection on page 382](#)
- [source \(PIM RPF Selection\) on page 383](#)
- [spt-threshold on page 384](#)
- [static \(Protocols PIM\) on page 385](#)

- [threshold \(PIM BFD Detection Time\)](#) on page 386
- [threshold \(PIM BFD Transmit Interval\)](#) on page 387
- [transmit-interval \(PIM BFD Liveness Detection\)](#) on page 388
- [traceoptions \(Protocols PIM\)](#) on page 389
- [version \(BFD\)](#) on page 392
- [version \(PIM\)](#) on page 393
- [wildcard-source \(PIM RPF Selection\)](#) on page 394

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 pim rp local (inet inet6) anycast-pim rp-set],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp local (inet inet6) anycast-pim rp-set],</p> <p>[edit protocols pim rp local (inet inet6) anycast-pim rp-set],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp local (inet inet6) anycast-pim rp-set]</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>
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 pim rp local family (inet inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6)], [edit protocols pim rp local family (inet inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. 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">• Configuring Local PIM RPs on page 181

address (Static RPs)

Syntax	<pre>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 static],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp static],</p> <p>[edit protocols pim static],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp static]</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>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure static rendezvous point (RP) addresses. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.</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>
Options	<p>address—Static RP address.</p> <p>Default: 224.0.0.0/4</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 the Static PIM RP Address on the Non-RP Routing Device on page 183

algorithm

Syntax	<code>algorithm <i>algorithm-name</i>;</code>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the algorithm to use for BFD authentication.
Options	<p><i>algorithm-name</i>—Name of algorithm to use for BFD authentication:</p> <ul style="list-style-type: none">• simple-password—Plain-text password. One to 16 bytes of plain text. One or more passwords can be configured.• keyed-md5—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive rates greater than 100 ms.• meticulous-keyed-md5—Meticulous keyed Message Digest 5 hash algorithm.• keyed-sha-1—Keyed Secure Hash Algorithm I for sessions with transmit and receive rates greater than 100 ms.• meticulous-keyed-sha-1—Meticulous keyed Secure Hash Algorithm I.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Bidirectional Forwarding Detection Authentication for PIM• Configuring BFD Authentication for PIM on page 135• authentication (Protocols PIM) on page 329

anycast-pim

Syntax	<pre>anycast-pim { rp-set { address address <forward-msdp-sa>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols pim rp local family (inet inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6)], [edit protocols pim rp local family (inet inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6)]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure properties for anycast RP using PIM. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PIM Anycast With or Without MSDP on page 186

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.
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">• Example: Configuring the PIM Assert Timeout on page 222

authentication

Syntax	<pre>authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check; }</pre>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the algorithm, security keychain, and level of authentication for BFD sessions running on PIM interfaces.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BFD Authentication for PIM on page 135 • Configuring BFD for PIM on page 134 • Understanding Bidirectional Forwarding Detection Authentication for PIM • bfd-liveness-detection (Protocols PIM) on page 331 • key-chain (Protocols PIM) on page 355 • loose-check on page 358

auto-rp

Syntax	<pre>auto-rp { (announce discovery mapping); (mapping-agent-election no-mapping-agent-election); }</pre>
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 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
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	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 Auto-RP on page 195

bfd-liveness-detection

Syntax	<pre> bfd-liveness-detection { authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check; } detection-time { threshold <i>milliseconds</i>; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } version (0 1 automatic); } </pre>
Hierarchy Level	<p>[edit protocols pim interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>authentication option introduced 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 12.1 for the QFX Series.</p>
Description	<p>Configure bidirectional forwarding detection (BFD) timers and authentication for 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"> • Configuring BFD for PIM on page 134 • Configuring BFD Authentication for PIM on page 135

bootstrap

Syntax	<pre>bootstrap { family (inet inet6) { export [<i>policy-names</i>]; import [<i>policy-names</i>]; priority <i>number</i>; } }</pre>
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.
Description	Configure parameters to control bootstrap routers and messages. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<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>
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 334

bootstrap-import

Syntax	<code>bootstrap-import [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols pim rp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> pim rp],</code> <code>[edit protocols pim rp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim rp]</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.
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 333

bootstrap-priority

Syntax	<code>bootstrap-priority <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>],</p> <p>[edit protocols pim <i>rp</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</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>
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	<p><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.</p> <p>Range: 0 through 255</p> <p>Default: 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>

dense-groups

Syntax	<code>dense-groups { <i>addresses</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.
Description	Configure which groups are operating in dense mode.
Options	<i>addresses</i> —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 on page 161

detection-time (BFD for PIM)

Syntax	<pre> detection-time { threshold milliseconds; } </pre>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the clear bfd adaptation command to return BFD interval timers to their configured values. The clear bfd adaptation command is hitless, meaning that the command does not affect traffic flow on the routing device.</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"> • Configuring BFD for PIM on page 134 • bfd-liveness-detection on page 331 • threshold on page 386

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>
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 on page 139 • family (Protocols PIM) on page 344

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.
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"> • Configuring PIM Designated Router Election on Point-to-Point Links on page 133

dr-register-policy

Syntax	dr-register-policy [<i>policy-names</i>];
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>], [edit protocols pim <i>rp</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>]
Release Information	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Apply one or more policies to control outgoing PIM register messages.
Options	<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"> • Configuring Register Message Filters on a PIM RP and DR on page 211 • rp-register-policy on page 380

embedded-rp

Syntax	<pre>embedded-rp { group-ranges { destination-ip-prefix </prefix-length>; } maximum-rps limit; }</pre>
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.
Description	Configure properties for embedded IP version 6 (IPv6) RPs. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring PIM Embedded RP for IPv6</i>

export (Bootstrap)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bootstrap family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap family (inet inet6)],</p> <p>[edit protocols pim rp bootstrap family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap family (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	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> • import (Protocols PIM Bootstrap) on page 350

export (Protocols PIM)

Syntax	<code>export [<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 9.6. 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 join and prune messages. PIM join and prune filters can be applied to PIM-SM and PIM-SSM messages. PIM join and prune filters cannot be applied to PIM-DM messages.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Filtering Outgoing PIM Join Messages on page 208

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>
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 (Protocols PIM)

Syntax	family (inet inet6) { disable; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols pim], [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], [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], [edit protocols pim interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Disable the PIM protocol for the specified family.
Options	inet —Disable the PIM protocol for the IP version 4 (IPv4) address family. inet6 —Disable the PIM protocol for the IP version 6 (IPv6) address family.
Related Documentation	<ul style="list-style-type: none">• Disabling PIM on page 139• <i>disable (PIM Graceful Restart)</i>• disable (PIM) on page 338

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>
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 on page 181

group (RPF Selection)

Syntax	<pre>group group-address{ source source-address { next-hop next-hop-address; } wildcard-source { next-hop next-hop-address; } }</pre>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> edit protocols pim rpf-selection]
Release Information	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the PIM group address for which you configure RPF selection group (RPF Selection) .
Default	By default, PIM RPF selection is not configured.
Options	group-address —PIM group address for which you configure RPF selection.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring PIM RPF Selection</i>

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"> • Configuring Local PIM RPs on page 181 • Configuring PIM Embedded RP for IPv6 • Example: Configuring Bidirectional PIM

hello-interval

Syntax	<code>hello-interval <i>seconds</i>;</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. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Specify how often the routing device sends PIM hello packets out of an interface.
Options	<i>seconds</i> —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 (Protocols PIM) on page 349• Modifying the PIM Hello Interval on page 128

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>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</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"> • Configuring Local PIM RPs on page 181 • <i>Example: Configuring Bidirectional PIM</i>

import (Protocols PIM Bootstrap)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols pim rp bootstrap (inet inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap (inet inet6)], [edit protocols pim rp bootstrap (inet inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap (inet inet6)]
Release Information	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. 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 (Bootstrap) on page 341

import (Protocols PIM)

Syntax	<code>import [<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 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 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	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"> • Filtering Incoming PIM Join Messages on page 209

infinity

Syntax	<code>infinity [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols pim spt-threshold],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>pim spt-threshold],</code> <code>[edit protocols pim spt-threshold],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim spt-threshold]</code>
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the infinity statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair.
Options	<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 on page 224

interface

Syntax	<pre> interface (all <i>interface-name</i>) { disable; family (inet inet6) { disable; } hello-interval <i>seconds</i>; mode (dense sparse sparse-dense); neighbor-policy [<i>policy-names</i>]; override-interval <i>milliseconds</i>; priority <i>number</i>; propagation-delay <i>milliseconds</i>; reset-tracking-bit; version <i>version</i>; } </pre>
Hierarchy Level	[edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
Release Information	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 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	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"> • PIM on Aggregated Interfaces on page 128

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.
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">• Configuring PIM Join Load Balancing on page 147• <code>clear pim join-distribution</code> in the CLI Explorer

join-prune-timeout

Syntax	<code>join-prune-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 in Junos OS Release 8.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the timeout for the join state. If the periodic join refresh message is not received before the timeout expires, the join state is removed.
Options	seconds —Number of seconds to wait for the periodic join message to arrive. Range: 210 through 240 seconds Default: 210 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Modifying the Join State Timeout on page 151

key-chain

Syntax	<code>key-chain <i>key-chain-name</i>;</code>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the security keychain to use for BFD authentication.
Options	key-chain-name —Name of the security keychain to use for BFD authentication. The name is a unique integer between 0 and 63. This must match one of the keychains in the authentication-key-chains statement at the [edit security] hierarchy level.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BFD Authentication for PIM on page 135 • Understanding Bidirectional Forwarding Detection Authentication for PIM • authentication (Protocols PIM) on page 329

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>The remaining statements are explained separately.</p>
Description	Configure the routing device's RP properties.
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 on page 181

local-address (Protocols PIM)

Syntax	<code>local-address <i>address</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols pim rp local family (inet inet6) anycast-pim],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6) anycast-pim],</p> <p>[edit protocols pim rp local family (inet inet6) anycast-pim],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp local family (inet inet6) anycast-pim]</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>
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	<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 PIM Anycast With or Without MSDP on page 186

loose-check

Syntax	loose-check;
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.</p> <p>By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure <i>loose checking</i>. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.</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">• Configuring BFD Authentication for PIM on page 135• Understanding Bidirectional Forwarding Detection Authentication for PIM• authentication (Protocols PIM) on page 329

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.
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"> • Configuring PIM Auto-RP on page 195

maximum-rps

Syntax	<code>maximum-rps <i>limit</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols pim <i>rp embedded-rp</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>pim <i>rp embedded-rp</i>],</code> <code>[edit protocols pim <i>rp embedded-rp</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp embedded-rp</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.
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	<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>Configuring PIM Embedded RP for IPv6</i>

minimum-interval (PIM BFD Liveness Detection)

Syntax	<code>minimum-interval <i>milliseconds</i>;</code>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the transmit-interval minimum-interval and minimum-receive-interval statements.
Options	<i>milliseconds</i> —Minimum transmit and receive interval. Range: 1 through 255,000 milliseconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BFD for PIM on page 134

minimum-interval (PIM BFD Transmit Interval)

Syntax	<code>minimum-interval <i>milliseconds</i>;</code>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the minimum interval after which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the minimum-interval statement at the [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection] hierarchy level.
Options	<i>milliseconds</i> —Minimum transmit interval value. Range: 1 through 255,000



NOTE: The threshold value specified in the **threshold** statement must be greater than the value specified in the **minimum-interval** statement for the **transmit-interval** statement.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BFD for PIM on page 134• bfd-liveness-detection on page 331• minimum-interval on page 361• threshold on page 387

minimum-receive-interval

Syntax	<code>minimum-receive-interval <i>milliseconds</i>;</code>
Hierarchy Level	<code>[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]</code>
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the minimum-interval statement at the <code>[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection]</code> hierarchy level.
Options	<i>milliseconds</i> —Minimum receive interval. Range: 1 through 255,000 milliseconds
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BFD for PIM on page 134

mode (Protocols PIM)

Syntax	<code>mode (dense sparse sparse-dense);</code>
Hierarchy Level	<code>[edit protocols pim interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>]</code>
Release Information	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 PIM to operate in sparse, dense, or sparse-dense mode.
Options	dense —Operate in dense mode. sparse —Operate in sparse mode. sparse-dense —Operate in sparse-dense mode. Default: sparse
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.

multiplier

Syntax	<code>multiplier <i>number</i>;</code>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.
Options	<i>number</i> —Number of hello packets. Range: 1 through 255 Default: 3
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BFD for PIM on page 134

neighbor-policy

Syntax	<code>neighbor-policy [<i>policy-names</i>];</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 in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Apply a PIM interface-level policy to filter neighbor IP addresses.
Options	<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">• Configuring Interface-Level PIM Neighbor Policies on page 207

next-hop (PIM RPF Selection)

Syntax	<code>next-hop <i>next-hop-address</i>;</code>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source]
Release Information	Statement introduced in JUNOS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the specific next-hop address for the PIM group source.
Options	<i>next-hop-address</i> —Specific next-hop address for the PIM group source.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PIM RPF Selection

no-adaptation (PIM BFD Liveness Detection)

Syntax	<code>no-adaptation;</code>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 9.0 Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure BFD sessions not to adapt to changing network conditions. We recommend that you <i>do not</i> disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BFD for PIM on page 134 • bfd-liveness-detection on page 331

override-interval

Syntax	override-interval <i>milliseconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols pim], [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], [edit protocols pim interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim] [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 10.1. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Set the maximum time in milliseconds to delay sending override join messages for a multicast network that has join suppression enabled. When a router or switch sees a prune message for a join it is currently suppressing, it waits for the interval specified by the override timer before it sends an override join message.
Options	This is a random timer with a value in milliseconds. Range: 0 through maximum override value Default: 2000 milliseconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Enabling Join Suppression on page 151• propagation-delay on page 374• reset-tracking-bit on page 376

pim

```

Syntax  pim {
    disable;
    assert-timeout seconds;
    dense-groups {
        addresses;
    }
    dr-election-on-p2p;
    export;
    family (inet | inet6) {
        disable;
    }
    graceful-restart {
        disable;
        restart-duration seconds;
    }
    import [ policy-names ];
    interface interface-name {
        accept-remote-source;
        disable;
        family (inet | inet6) {
            disable;
        }
        hello-interval seconds;
        mode (dense | sparse | sparse-dense);
        neighbor-policy [ policy-names ];
        override-interval milliseconds;
        priority number;
        propagation-delay milliseconds;
        reset-tracking-bit;
        version version;
    }
    join-load-balance;
    join-prune-timeout;
    nonstop-routing;
    override-interval milliseconds;
    propagation-delay milliseconds;
    reset-tracking-bit;
    rib-group group-name;
    rp {
        auto-rp {
            (announce | discovery | mapping);
            (mapping-agent-election | no-mapping-agent-election);
        }
        bootstrap {
            family (inet | inet6) {
                export [ policy-names ];
                import [ policy-names ];
                priority number;
            }
        }
        bootstrap-import [ policy-names ];
        bootstrap-export [ policy-names ];
    }
}

```

```

bootstrap-priority number;
dr-register-policy [ policy-names ];
embedded-rp {
    group-ranges {
        destination-ip-prefix </prefix-length>;
    }
    maximum-rps limit;
}
local {
    family (inet | inet6) {
        address address;
        anycast-pim {
            disable;
            rp-set {
                address address <forward-msdp-sa>;
            }
            local-address address;
        }
        group-ranges {
            destination-ip-prefix </prefix-length>;
        }
        hold-time seconds;
        priority number;
    }
}
rp-register-policy [ policy-names ];
spt-threshold {
    infinity [ policy-names ];
}
static {
    address address {
        group-ranges {
            version version;
            destination-ip-prefix </prefix-length>;
        }
    }
}
rpf-selection {
    group group-address {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
    prefix-list prefix-list-addresses {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
}
traceoptions {

```

```

    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  tunnel-devices [ mt-fpc/pic/port ];
}

```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. family statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Enable PIM on the routing device. The statements are explained separately.
Default	PIM is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

prefix-list (PIM RPF Selection)

Syntax	<pre>prefix-list <i>prefix-list-addresses</i> { source <i>source-address</i> { next-hop <i>next-hop-address</i>; } wildcard-source { next-hop <i>next-hop-address</i>; } }</pre>
Hierarchy Level	<pre>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source]</pre>
Release Information	Statement introduced in Junos OS Release 10.4. 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	(Optional) Configure a list of prefixes (addresses) for multiple PIM groups.
Options	<p><i>prefix-list-addresses</i>—List of prefixes (addresses) for multiple PIM groups.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring PIM RPF Selection</i>

priority (Bootstrap)

Syntax	<code>priority <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bootstrap (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap (inet inet6)],</p> <p>[edit protocols pim rp bootstrap (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap (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>number—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 335

priority (PIM Interfaces)

Syntax	<code>priority <i>number</i>;</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. 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	<i>number</i> —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">• Configuring Interface Priority for PIM Designated Router Selection on page 132

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"> • Configuring Local PIM RPs on page 181 • Example: Configuring Bidirectional PIM

propagation-delay

Syntax	<code>propagation-delay <i>milliseconds</i>;</code>
Hierarchy Level	[edit protocols pim], [edit protocols pim interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols pim], [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>]
Release Information	Statement introduced in Junos OS Release 10.1. 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	Set a delay for implementing a PIM prune message on the upstream routing device on a multicast network for which join suppression has been enabled. The routing device waits for the prune pending period to detect whether a join message is currently being suppressed by another routing device.
Options	<i>milliseconds</i> —Interval for the prune pending timer, which is the sum of the propagation-delay value and the override-interval value. Range: 250 through 2000 milliseconds Default: 500 milliseconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Enabling Join Suppression on page 151• override-interval on page 366• reset-tracking-bit on page 376

register-probe-time

Syntax	<code>register-probe-time</code> <i>register-probe-time</i> ;
Hierarchy Level	[edit protocols pim rp]
Release Information	Statement introduced in Junos OS Release 12.2 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D16 for QFX Series switches.
Description	Specify the amount of time before the register suppression time (RST) expires when a designated switch can send a NULL-Register to the rendezvous point (RP).
Options	<i>register-probe-time</i> —Amount of time before the RST expires. Default: 5 seconds Range: 5 to 60 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">• PIM Overview on page 125• Understanding PIM Sparse Mode on page 143

reset-tracking-bit

Syntax	reset-tracking-bit;
Hierarchy Level	[edit protocols pim], [edit protocols pim interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols pim], [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>]
Release Information	Statement introduced in Junos OS Release 10.1. 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	Change the value of a tracking bit (T-bit) field in the LAN prune delay hello option from the default of 1 to 0, which enables join suppression for a multicast interface. When the network starts receiving multiple identical join messages, join suppression triggers a random timer with a value of 66 through 84 milliseconds ($1.1 \times \text{periodic}$ through $1.4 \times \text{periodic}$, where periodic is 60 seconds). This creates an interval during which no identical join messages are sent. Eventually, only one of the identical messages is sent. Join suppression is triggered each time identical messages are sent for the same join.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Enabling Join Suppression on page 151• override-interval on page 366• propagation-delay on page 374

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>

rp

```

Syntax  register-probe-time {
        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 on page 143](#)

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 on page 211](#)
- [dr-register-policy on page 339](#)

rp-set

Syntax	<pre>rp-set { address address <forward-msdp-sa>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols pim local family (inet inet6) anycast-pim],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim local family (inet inet6) anycast-pim],</p> <p>[edit protocols pim local family (inet inet6) anycast-pim],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim local family (inet inet6) anycast-pim]</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 a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 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"> • Example: Configuring PIM Anycast With or Without MSDP on page 186

rpf-selection

Syntax	<pre> 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; } } } </pre>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols pim]
Release Information	<p>Statement introduced in JUNOS Release 10.4.</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 PIM RPF next-hop neighbor for a specific group and source for a VRF routing instance.</p> <p>The remaining statements are explained separately.</p>
Default	If you omit the rpf-selection statement, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows.
Options	source-address —Specific source address for the PIM group.
Required Privilege Level	<p>view-level—To view this statement in the configuration.</p> <p>control-level—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Example: Configuring PIM RPF Selection</i>

source (PIM RPF Selection)

Syntax	<pre>source source-address { next-hop next-hop-address; }</pre>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i>]
Release Information	Statement introduced in JUNOS Release 10.4. 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 source address for the PIM group.
Options	<p>source-address—Specific source address for the PIM group.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring PIM RPF Selection</i>


spt-threshold

Syntax	spt-threshold { infinity [<i>policy-names</i>]; }
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">• Example: Configuring the PIM SPT Threshold Policy on page 224


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],</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 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.</p> <p>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 on page 183

threshold (PIM BFD Detection Time)

Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection detection-time], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection detection-time]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.
<div> NOTE: The threshold value must be equal to or greater than the transmit interval.</div> <div>The threshold time must be equal to or greater than the value specified in the minimum-interval or the minimum-receive-interval statement.</div>	
Options	<i>milliseconds</i> —Value for the detection time adaptation threshold. Range: 1 through 255,000
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BFD for PIM on page 134• bfd-liveness-detection on page 331• detection-time on page 337• minimum-interval on page 361• minimum-receive-interval on page 363

threshold (PIM BFD Transmit Interval)

Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.
Options	<i>milliseconds</i> —Value for the transmit interval adaptation threshold. Range: 0 through 4,294,967,295 ($2^{32} - 1$)
<div>  <p>NOTE: The threshold value specified in the <code>threshold</code> statement must be greater than the value specified in the <code>minimum-interval</code> statement for the <code>transmit-interval</code> statement.</p> </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 BFD for PIM on page 134 • bfd-liveness-detection on page 331

transmit-interval (PIM BFD Liveness Detection)

Syntax	<pre>transmit-interval { minimum-interval milliseconds; threshold milliseconds; }</pre>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify the transmit interval for the bfd-liveness-detection statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.</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">• Configuring BFD for PIM on page 134• bfd-liveness-detection on page 331• threshold on page 387• minimum-interval on page 362• minimum-receive-interval on page 363

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

Related Documentation	<ul style="list-style-type: none"> • Configuring PIM Trace Options on page 130 • Tracing DVMRP Protocol Traffic • Tracing MSDP Protocol Traffic on page 234 • Configuring PIM Trace Options on page 130
------------------------------	---

version (BFD)

Syntax	version (0 1 automatic);
Hierarchy Level	[edit protocols piminterface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the bidirectional forwarding detection (BFD) protocol version that you want to detect.
Options	Configure the BFD version to detect: 1 (BFD version 1) or automatic (autodetect the BFD version) Default: automatic
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BFD for PIM on page 134

version (PIM)

Syntax	<code>version <i>version</i>;</code>
Hierarchy Level	<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 static address <i>address</i>],</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 static address <i>address</i>],</p> <p>[edit protocols pim interface <i>interface-name</i>],</p> <p>[edit protocols pim rp static address <i>address</i>],</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 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>
Description	Specify the version of PIM.
Options	<p>version—PIM version number.</p> <p>Range: 1 or 2</p> <p>Default: PIMv1 for rendezvous point (RP) mode (at the [edit protocols pim rp static address <i>address</i>] hierarchy level). PIMv2 for interface mode (at the [edit protocols pim interface <i>interface-name</i>] hierarchy level).</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 PIM Sparse Mode on page 146 • Configuring PIM Dense Mode Properties on page 160 • Configuring PIM Sparse-Dense Mode Properties on page 161

wildcard-source (PIM RPF Selection)

Syntax	wildcard-source { next-hop next-hop-address; }
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i>]
Release Information	Statement introduced in Junos OS Release 10.4. 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	Use a wildcard for the multicast source instead of (or in addition to) a specific multicast source. The remaining statements are explained separately.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring PIM RPF Selection</i>

CHAPTER 21

Configuration Statements (Source-Specific Multicast)

- [asm-override-ssm on page 395](#)
- [policy \(SSM Maps\) on page 396](#)
- [ssm-groups on page 397](#)
- [ssm-map \(Protocols IGMP\) on page 398](#)
- [ssm-map \(Routing Options Multicast\) on page 399](#)
- [ssm-map-policy \(IGMP\) on page 400](#)

asm-override-ssm

Syntax	asm-override-ssm;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Enable the routing device to accept any-source multicast join messages (*G) for group addresses that are within the default or configured range of source-specific multicast groups.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	• Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 172

policy (SSM Maps)

Syntax	<code>policy [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>],</code> <code>[edit routing-options multicast ssm-map <i>ssm-map-name</i>]</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 12.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Apply one or more policies to an SSM map.
Options	<i>policy-names</i> —Name of one or more policies for SSM mapping.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To view this statement in the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring SSM Mapping on page 169

ssm-groups

Syntax	<code>ssm-groups [<i>ip-addresses</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</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 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure source-specific multicast (SSM) groups.</p> <p>By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the ssm-groups statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the ssm-groups statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.</p> <p>IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.</p>
Options	<i>ip-addresses</i> —List of one or more additional SSM group addresses separated by a space.
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 Source-Specific Multicast Groups with Any-Source Override on page 172

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 <i>interface-name</i>], [edit protocols igmp interface <i>interface-name</i>]
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 on page 169

ssm-map (Routing Options Multicast)

Syntax	<pre>ssm-map <i>ssm-map-name</i> { <i>policy</i> [<i>policy-names</i>]; source [<i>addresses</i>]; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</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 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure SSM mapping.
Options	<p><i>ssm-map-name</i>—Name of the SSM map.</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 SSM Mapping on page 169

ssm-map-policy (IGMP)

Syntax	<code>ssm-map-policy <i>ssm-map-policy-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 11.4. 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 policy to an IGMP interface.
Options	<i>ssm-map-policy-name</i> —Name of SSM map policy.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring SSM Maps for Different Groups to Different Sources on page 175

CHAPTER 22

Configuration Statements (MSDP)

- [active-source-limit](#) on page 402
- [authentication-key](#) on page 403
- [data-encapsulation](#) on page 404
- [default-peer](#) on page 405
- [disable \(Protocols MSDP\)](#) on page 406
- [export \(Protocols MSDP\)](#) on page 407
- [group \(Protocols MSDP\)](#) on page 408
- [import \(Protocols MSDP\)](#) on page 409
- [local-address \(Protocols MSDP\)](#) on page 410
- [maximum \(MSDP Active Source Messages\)](#) on page 411
- [mode \(Protocols MSDP\)](#) on page 412
- [msdp](#) on page 413
- [peer \(Protocols MSDP\)](#) on page 415
- [rib-group \(Protocols MSDP\)](#) on page 416
- [source \(Protocols MSDP\)](#) on page 417
- [threshold \(MSDP Active Source Messages\)](#) on page 418
- [traceoptions \(Protocols MSDP\)](#) on page 419

active-source-limit

Syntax	<pre>active-source-limit { log-interval <i>seconds</i>; log-warning <i>value</i>; maximum <i>number</i>; threshold <i>number</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols msdp], [edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols msdp peer <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols msdp source <i>ip-address/prefix-length</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols msdp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp source <i>ip-address/prefix-length</i>], [edit protocols msdp], [edit protocols msdp group <i>group-name</i> peer <i>address</i>], [edit protocols msdp peer <i>address</i>], [edit protocols msdp source <i>ip-address/prefix-length</i>], [edit routing-instances <i>routing-instance-name</i> protocols msdp], [edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols msdp source <i>ip-address/prefix-length</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	Limit the number of active source messages the routing device accepts.
Default	If you do not include this statement, the router accepts any number of MSDP active source messages.
Options	The options are explained separately.
Required Privilege Level	<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 MSDP with Active Source Limits and Mesh Groups on page 238

authentication-key

Syntax	<code>authentication-key peer-key;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <code>msdp peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>msdp peer address</code>],</p> <p>[edit protocols <code>msdp group group-name peer address</code>],</p> <p>[edit protocols <code>msdp peer address</code>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>msdp peer address</code>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	Associate a Message Digest 5 (MD5) signature option authentication key with an MSDP peering session.
Default	If you do not include this statement, the routing device accepts any valid MSDP messages from the peer address.
Options	<p>peer-key—MD5 authentication key. The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of (,), &, and [. If you include spaces in an MSDP authentication key, enclose all characters in quotation marks (" ").</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 MSDP in a Routing Instance</i>

data-encapsulation

Syntax	<code>data-encapsulation (disable enable);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols msdp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp], [edit protocols msdp], [edit routing-instances <i>routing-instance-name</i> protocols msdp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a rendezvous point (RP) using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages.
Default	If you do not include this statement, the RP encapsulates multicast data.
Options	disable —(Optional) Do not use MSDP data encapsulation. enable —Use MSDP data encapsulation. Default: <code>enable</code>
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">• Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 238

default-peer

Syntax	default-peer;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit protocols msdp],</p> <p>[edit protocols msdp group <i>group-name</i>],</p> <p>[edit protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit protocols msdp peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	Establish this peer as the default MSDP peer and accept source-active messages from the peer without the usual peer-reverse-path-forwarding (peer-RPF) check.
Required Privilege Level	<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 MSDP with Active Source Limits and Mesh Groups on page 238

disable (Protocols MSDP)

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit protocols msdp],</p> <p>[edit protocols msdp group <i>group-name</i>],</p> <p>[edit protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit protocols msdp peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	Explicitly disable MSDP.
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>Disabling MSDP</i>

export (Protocols MSDP)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit protocols msdp],</p> <p>[edit protocols msdp group <i>group-name</i>],</p> <p>[edit protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit protocols msdp peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	Apply one or more policies to routes being exported from the routing table into MSDP.
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>Example: Configuring MSDP in a Routing Instance</i> • import on page 409

group (Protocols MSDP)

Syntax	<pre> group <i>group-name</i> { disable; export [<i>policy-names</i>]; import [<i>policy-names</i>]; local-address <i>address</i>; mode (mesh-group standard); traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; } peer <i>address</i>; { disable; active-source-limit { maximum <i>number</i>; threshold <i>number</i>; } authentication-key <i>peer-key</i>; default-peer; export [<i>policy-names</i>]; import [<i>policy-names</i>]; local-address <i>address</i>; traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; } } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit protocols msdp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	<p>Define an MSDP peer group. MSDP peers within groups share common tracing options, if present and not overridden for an individual peer with the peer statement. To configure multiple MSDP groups, include multiple group statements.</p> <p>By default, the group's options are identical to the global MSDP options. To override the global options, include group-specific options within the group statement.</p> <p>The group must contain at least one peer.</p>
Options	<p>group-name—Name of the MSDP group.</p> <p>The remaining statements are explained separately.</p>

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"> • <i>Example: Configuring MSDP in a Routing Instance</i>

import (Protocols MSDP)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit protocols msdp],</p> <p>[edit protocols msdp group <i>group-name</i>],</p> <p>[edit protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit protocols msdp peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	Apply one or more policies to routes being imported into the routing table from MSDP.
Options	<i>policy-names</i> —Name of one or more policies.
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"> • <i>Example: Configuring MSDP in a Routing Instance</i> • export (Protocols MSDP) on page 407

local-address (Protocols MSDP)

Syntax	<code>local-address address;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit protocols msdp],</p> <p>[edit protocols msdp group <i>group-name</i>],</p> <p>[edit protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit protocols msdp peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	Configure the local end of an MSDP session. You must configure at least one peer for MSDP to function. When configuring a peer, you must include this statement. This address is used to accept incoming connections to the peer and to establish connections to the remote peer.
Options	address —IP address of the local end of the connection.
Required Privilege Level	<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 MSDP in a Routing Instance</i>

maximum (MSDP Active Source Messages)

Syntax	<code>maximum <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols msdp active-source-limit],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp active-source-limit],</p> <p>[edit protocols msdp active-source-limit],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp active-source-limit]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	Configure the maximum number of MSDP active source messages the router accepts.
Options	<p><i>number</i>—Maximum number of active source messages.</p> <p>Range: 1 through 1,000,000</p> <p>Default: 25,000</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 MSDP with Active Source Limits and Mesh Groups on page 238 • threshold (MSDP Active Source Messages) on page 418

mode (Protocols MSDP)

Syntax	mode (mesh-group standard);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>], [edit protocols msdp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure groups of peers in a full mesh topology to limit excessive flooding of source-active messages to neighboring peers. The default flooding mode is standard .
Default	If you do not include this statement, default flooding is applied.
Options	mesh-group —Group of peers that are mesh group members. standard —Use standard MSDP source-active flooding rules. Default: standard
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 238

msdp

```

Syntax  msdp {
        disable;
        active-source-limit {
            log-interval seconds;
            log-warning value;
            maximum number;
            threshold number;
        }
        data-encapsulation (disable | enable);
        export [ policy-names ];
        group group-name {
            ...group-configuration ...
        }
        hold-time seconds;
        import [ policy-names ];
        local-address address;
        keep-alive seconds;
        peer address {
            ...peer-configuration ...
        }
        rib-group group-name;
        source ip-prefix</prefix-length> {
            active-source-limit {
                maximum number;
                threshold number;
            }
        }
        sa-hold-time seconds;
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
        group group-name {
            disable;
            export [ policy-names ];
            import [ policy-names ];
            local-address address;
            mode (mesh-group | standard);
            peer address {
                ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
                just following ...
            }
            traceoptions {
                file filename <files number> <size size> <world-readable | no-world-readable>;
                flag flag <flag-modifier> <disable>;
            }
        }
        peer address {
            disable;
            active-source-limit {
                maximum number;
                threshold number;
            }
        }
    }

```

```
    }  
    authentication-key peer-key;  
    default-peer;  
    export [ policy-names ];  
    import [ policy-names ];  
    local-address address;  
    traceoptions {  
        file filename <files number> <size size> <world-readable | no-world-readable>;  
        flag flag <flag-modifier> <disable>;  
    }  
}  
}
```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable MSDP on the router or switch. You must also configure at least one peer for MSDP to function.
Default	MSDP is disabled on the router or switch.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring MSDP in a Routing Instance</i>

peer (Protocols MSDP)

Syntax	<pre> peer address { disable; active-source-limit { maximum number; threshold number; } authentication-key peer-key; default-peer; export [policy-names]; import [policy-names]; local-address address; traceoptions { file filename <files number> <size size> <world-readable no-world-readable>; flag flag <flag-modifier> <disable>; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit protocols msdp],</p> <p>[edit protocols msdp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	<p>Define an MSDP peering relationship. An MSDP routing device must know which routing devices are its peers. You define the peer relationships explicitly by configuring the neighboring routing devices that are the MSDP peers of the local routing device. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. To configure multiple MSDP peers, include multiple peer statements.</p> <p>By default, the peer's options are identical to the global or group-level MSDP options. To override the global or group-level options, include peer-specific options within the peer (Protocols MSDP) statement.</p> <p>At least one peer must be configured for MSDP to function. You must configure address and local-address.</p>
Options	<p>address—Name of the MSDP peer.</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

- *Example: Configuring MSDP in a Routing Instance*

rib-group (Protocols MSDP)

Syntax `rib-group group-name;`

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

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

Description Associate a routing table group with MSDP.

Options *group-name*—Name of the routing table group. The name must be one that you defined with the **rib-groups** statement at the [edit routing-options] hierarchy level.

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

Related Documentation

- *Example: Configuring MSDP in a Routing Instance*

source (Protocols MSDP)

Syntax	<pre>source ip-address </prefix-length> { active-source-limit { maximum number; threshold number; } }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit protocols msdp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	Limit the number of active source messages the routing device accepts from sources in this address range.
Default	If you do not include this statement, the routing device accepts any number of MSDP active source messages.
Options	The other statements are explained separately.
Required Privilege Level	<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 MSDP with Active Source Limits and Mesh Groups on page 238

threshold (MSDP Active Source Messages)

Syntax	<code>threshold <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols msdp active-source-limit], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp active-source-limit], [edit protocols msdp active-source-limit], [edit routing-instances <i>routing-instance-name</i> protocols msdp active-source-limit]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the random early detection (RED) threshold for MSDP active source messages. This number must be less than the configured or default maximum.
Options	<i>number</i> —RED threshold for active source messages. Range: 1 through 1,000,000 Default: 24,000
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 238• maximum (MSDP Active Source Messages) on page 411

traceoptions (Protocols MSDP)

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 msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit protocols msdp],</p> <p>[edit protocols msdp group <i>group-name</i>],</p> <p>[edit protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit protocols msdp peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</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	<p>Configure MSDP tracing options.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	<p>The default MSDP trace options are those inherited from the routing protocol's traceoptions statement included at the [edit routing-options] hierarchy level.</p>
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 msdp-log file.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p>

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

Range: 2 through 1000 files

Default: 2 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

MSDP Tracing Flags

- **keepalive**—Keepalive messages
- **packets**—All MSDP packets
- **route**—MSDP changes to the routing table
- **source-active**—Source-active packets
- **source-active-request**—Source-active request packets
- **source-active-response**—Source-active response packets

Global Tracing Flags

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

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

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

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

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

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

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

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

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

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

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

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

Syntax: **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 Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Tracing MSDP Protocol Traffic on page 234

CHAPTER 23

Operational Commands (IGMP)

- `clear igmp membership`
- `clear igmp statistics`
- `show igmp group`
- `show configuration protocols igmp`
- `show igmp interface`
- `show igmp statistics`
- `show system statistics igmp`

clear igmp membership

List of Syntax	Syntax on page 424 Syntax (EX Series Switch and the QFX Series) on page 424
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	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) 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 429• show igmp interface on page 435
List of Sample Output	clear igmp membership on page 424 clear igmp membership interface on page 425 clear igmp membership group on page 426
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 427 Syntax (EX Series Switches) on page 427
Syntax	clear igmp statistics <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	clear igmp statistics <interface <i>interface-name</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.
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
List of Sample Output	clear igmp statistics on page 427
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            0        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		

show igmp group

List of Syntax	Syntax on page 429 Syntax (EX Series Switch and the QFX Series) on page 429
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>
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
List of Sample Output	show igmp group (Include Mode) on page 430 show igmp group (Exclude Mode) on page 431 show igmp group brief on page 431 show igmp group detail on page 431
Output Fields	<p>Table 12 on page 429 describes the output fields for the show igmp group command. Output fields are listed in the approximate order in which they appear.</p>

Table 12: show igmp group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the IGMP membership report. A name of local indicates that the local routing device joined the group itself.	All levels
Group	Group address.	All levels
Group Mode	Mode the SSM group is operating in: Include or Exclude .	All levels
Source	Source address.	All levels

Table 12: show igmp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 configuration protocols igmp

Syntax	show configuration protocols igmp
Release Information	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 Internet Group Management Protocol (IGMP) information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • IGMP Snooping Overview on page 65 • Configuring IGMP Snooping on page 68
List of Sample Output	show configuration protocols igmp on page 433
Output Fields	Table 12 on page 429 describes the output fields for the show configuration protocols igmp command that relate to IGMP querying.

Table 13: show igmp group Output Fields

Field Name	Field Description	Level of Output
accounting	Enables notification for join and leave events.	All levels
igmp-querier	Configured source address for the IGMP querier.	All levels
interface	Name of the interface that receives IGMP membership reports.	All levels
query-interval	Interval at which the IGMP querier sends general host-query messages to solicit membership information.	All levels
query-response-interval	How long the IGMP querier waits to receive a response from a query message before sending another query.	All levels
src-address	Source address of IGMP queries.	
version	IGMP version.	All levels

Sample Output

show configuration protocols igmp

```

user@switch> show configuration protocols igmp
query-interval 150;
query-response-interval 50;
accounting;
interface vlan.43 {
  version 2;
}
igmp-querier {

```

```
src-address 10.0.0.2;  
}
```

show igmp interface

List of Syntax	Syntax on page 435 Syntax (EX Series Switch and the QFX Series) on page 435
Syntax	<pre>show igmp interface <brief detail> <interface-name> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switch 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>
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 424
List of Sample Output	show igmp interface on page 437 show igmp interface brief on page 437 show igmp interface detail on page 438
Output Fields	<p>Table 14 on page 435 describes the output fields for the show igmp interface command. Output fields are listed in the approximate order in which they appear.</p>

Table 14: show igmp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the routing device that has been elected to send membership queries.	All levels

Table 14: show igmp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	State of the interface: Up or Down .	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.	All levels
Timeout	How long until the IGMP querier is declared to be unreachable, in seconds.	All levels
Version	IGMP version being used on the interface: 1 , 2 , or 3 .	All levels
Groups	Number of groups on the interface.	All levels
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
Passive	State of the passive mode option: <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

Table 14: show igmp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Configured Parameters	Information configured by the user: <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	Derived information: <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 437](#).

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 437](#).

show igmp statistics

List of Syntax	Syntax on page 439 Syntax (EX Series Switch and the QFX Series) on page 439
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>
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 427
List of Sample Output	show igmp statistics on page 440 show igmp statistics interface on page 441
Output Fields	<p>Table 15 on page 439 describes the output fields for the show igmp statistics command. Output fields are listed in the approximate order in which they appear.</p>

Table 15: 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 15: 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 system statistics igmp

List of Syntax	Syntax on page 442 Syntax (EX Series Switches) on page 442 Syntax (TX Matrix Router) on page 442 Syntax (TX Matrix Plus Router) on page 442
Syntax	show system statistics igmp
Syntax (EX Series Switches)	show system statistics igmp <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system statistics igmp <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system statistics igmp <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display system-wide Internet Group Management Protocol (IGMP) statistics.
Options	none —Display system statistics for IGMP. all-chassis —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system statistics for IGMP for all the routers in the chassis. all-lcc —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for IGMP for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for all connected T1600 or T4000 LCCs. all-members —(EX4200 switches only) (Optional) Display IGMP statistics for all members of the Virtual Chassis configuration. lcc <i>number</i> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for IGMP for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(EX4200 switches only) (Optional) Display IGMP statistics for the local Virtual Chassis member.

member *member-id*—(EX4200 switches only) (Optional) Display IGMP statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

scc—(TX Matrix routers only) (Optional) Display system statistics for IGMP for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display system statistics for IGMP for the TX Matrix Plus router. Replace *number* with 0.

Additional Information By default, when you issue the **show system statistics igmp** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level view

Related Documentation • [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [show system statistics igmp on page 443](#)
[show system statistics igmp \(EX Series Switches\) on page 444](#)
[show system statistics igmp \(TX Matrix Plus Router\) on page 444](#)

Sample Output

show system statistics igmp

```
user@host> show system statistics igmp
igmp:
    17178 messages received
    0 messages received with too few bytes
    0 messages received with bad checksum
    0 membership queries received
```

```
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent
```

show system statistics igmp (EX Series Switches)

```
user@host> show system statistics igmp
igmp:
    0 messages received
    0 messages received with too few bytes
    0 messages received with bad checksum
    0 membership queries received
    0 membership queries received with invalid fields
    0 membership reports received
    0 membership reports received with invalid fields
    0 membership reports received for groups to which we belong
    0 Membership reports sent
```

show system statistics igmp (TX Matrix Plus Router)

```
user@host> show system statistics igmp
sfc0-re0:
-----
igmp:
    0 messages received
    0 messages received with too few bytes
    0 messages received with bad checksum
    0 membership queries received
    0 membership queries received with invalid field(s)
    0 membership reports received
    0 membership reports received with invalid field(s)
    0 membership reports received for groups to which we belong
    0 membership reports sent
```

```
lcc0-re0:
-----
igmp:
    0 messages received
    0 messages received with too few bytes
    0 messages received with bad checksum
    0 membership queries received
    0 membership queries received with invalid field(s)
    0 membership reports received
    0 membership reports received with invalid field(s)
    0 membership reports received for groups to which we belong
    0 membership reports sent
```

```
lcc1-re0:
-----
igmp:
    0 messages received
    0 messages received with too few bytes
    0 messages received with bad checksum
    0 membership queries received
    0 membership queries received with invalid field(s)
    0 membership reports received
    0 membership reports received with invalid field(s)
    0 membership reports received for groups to which we belong
    0 membership reports sent
```


lcc2-re0:

igmp:

0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent

lcc3-re0:

igmp:

0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent

CHAPTER 24

Operational Commands (IGMP Snooping)

- clear igmp-snooping membership
- clear igmp-snooping statistics
- show igmp-snooping membership
- show igmp-snooping route
- show igmp-snooping statistics
- show igmp-snooping vlans

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 11.1 for the QFX Series.
Description	Clear IGMP snooping membership information.
Options	<code>vlan <i>vlan-name</i></code> —(Optional) Name of the VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show igmp-snooping membership on page 450
List of Sample Output	clear igmp-snooping membership on page 448

Sample Output

clear igmp-snooping membership

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

clear igmp-snooping statistics

Syntax	<code>clear igmp-snooping statistics</code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear IGMP snooping statistics.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show igmp-snooping statistics on page 455
List of Sample Output	clear igmp-snooping statistics on page 449

Sample Output

clear igmp-snooping statistics

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

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	<p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>IGMPv3 output introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display IGMP snooping membership information.
Options	<p>none—Display general parameters.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display IGMP snooping information for the specified interface.</p> <p>vlan <i>vlan-id</i> <i>vlan-name</i>—(Optional) Display IGMP snooping information for the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Monitoring IGMP Snooping on page 72 • Configuring IGMP Snooping on page 68 • show igmp-snooping route on page 453 • show igmp-snooping statistics on page 455 • show igmp-snooping vlans on page 457
List of Sample Output	<p>show igmp-snooping membership on page 451</p> <p>show igmp-snooping membership detail on page 452</p>
Output Fields	<p>Table 16 on page 450 lists the output fields for the show igmp-snooping membership command. Output fields are listed in the approximate order in which they appear.</p>

Table 16: show igmp-snooping membership Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All
Interfaces	Interfaces assigned to the VLAN.	All
Tag	Numerical identifier of the VLAN.	detail

Table 16: show igmp-snooping membership Output Fields (*continued*)

Field Name	Field Description	Level of Output
Router interfaces	Names of multicast router interfaces.	detail
• static or dynamic	Whether the multicast router interface is static or dynamic .	detail
• Uptime	For static interfaces, length of time since the interface was configured as a multicast router interface; for dynamic interfaces, length of time since the first query was received on the interface.	detail
• timeout	Query timeout in seconds.	detail
Group	IP multicast address of the multicast group.	detail
Receiver count	Number of interfaces that have membership in a multicast group.	detail
Flags	IGMP version of the host sending a join message.	detail
Uptime	Length of time a multicast group has been active on the interface.	detail
timeout	Time (in seconds) left until the entry for the multicast group is removed.	All
Last reporter	Last host to report membership for the multicast group.	detail
Include source	Source addresses from which multicast streams are allowed based on IGMPv3 reports.	detail

Sample Output

show igmp-snooping membership

```

user@switch> show igmp-snooping membership
VLAN: v1
  224.1.1.1      *           258 secs
    Interfaces: ge-0/0/0.0
  224.1.1.3      *           258 secs
    Interfaces: ge-0/0/0.0
  224.1.1.5      *           258 secs
    Interfaces: ge-0/0/0.0
  224.1.1.7      *           258 secs

```

```
Interfaces: ge-0/0/0.0
224.1.1.9      *           258 secs
Interfaces: ge-0/0/0.0
224.1.1.11     *           258 secs
Interfaces: ge-0/0/0.0
```

show igmp-snooping membership detail

```
user@switch> show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
Group: 225.0.0.2
Receiver count: 1, Flags: <V3-hosts>
  ge-0/0/15.0 Uptime: 00:00:11 timeout: 248 Last reporter: 10.2.10.16
  Include source: 1.2.1.1, 1.3.1.1
VLAN: v44 Tag: 44 (Index: 5)
Group: 225.0.0.1
Receiver count: 1, Flags: <V2-hosts>
  ge-0/0/21.0 Uptime: 00:00:02 timeout: 257
VLAN: v110 Tag: 110 (Index: 4)
Router interfaces:
  ge-0/0/3.0 static Uptime: 00:08:45
  ge-0/0/2.0 static Uptime: 00:08:45
  ge-0/0/4.0 dynamic Uptime: 00:16:41 timeout: 254
Group: 225.0.0.3
Receiver count: 1, Flags: <V3-hosts>
  ge-0/0/5.0 Uptime: 00:00:19 timeout: 259
Group: 225.1.1.1
Receiver count: 1, Flags: <V2-hosts>
  ge-0/0/5.0 Uptime: 00:22:43 timeout: 96
Group: 225.2.2.2
Receiver count: 1, Flags: <V2-hosts Static>
  ge-0/0/5.0 Uptime: 00:23:13
```


show igmp-snooping route

Syntax	<pre>show igmp-snooping route <brief detail> <ethernet-switching <brief detail vlan (vlan-id vlan-name)>> <inet <brief detail vlan vlan-name>> <vlan vlan-name></pre>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display IGMP snooping route information.
Options	<p>none—Display general parameters.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>ethernet-switching—(Optional) Display Ethernet switching information.</p> <p>inet—(Optional) Display inet information.</p> <p>vlan vlan-name—(Optional) Display route information for the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Monitoring IGMP Snooping on page 72 • Configuring IGMP Snooping on page 68 • show igmp-snooping statistics on page 455 • show igmp-snooping vlans on page 457
List of Sample Output	<p>show igmp-snooping route on page 454</p> <p>show igmp-snooping route vlan v1 on page 454</p>
Output Fields	<p>Table 17 on page 453 lists the output fields for the show igmp-snooping route command. Output fields are listed in the approximate order in which they appear.</p>

Table 17: show igmp-snooping route Output Fields

Field Name	Field Description
Table	(For internal use only. Value is always 0.)
VLAN	Name of the VLAN.
Group	Multicast group address.
Interfaces	Interfaces on which IGMP packets were snooped.
Next-hop	ID associated with the next-hop device.

Sample Output

show igmp-snooping route

```
user@switch> show igmp-snooping route
VLAN          Group          Next-hop
V11           224.1.1.1, *      533
               Interfaces: ge-0/0/13.0, ge-0/0/1.0
VLAN          Group          Next-hop
v12           224.1.1.3, *      534
               Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

show igmp-snooping route vlan v1

```
user@switch> show igmp-snooping route vlan v1
Table: 0
VLAN          Group          Next-hop
v1           224.1.1.1, *      1266
               Interfaces: ge-0/0/0.0
v1           224.1.1.3, *      1266
               Interfaces: ge-0/0/0.0
v1           224.1.1.5, *      1266
               Interfaces: ge-0/0/0.0
v1           224.1.1.7, *      1266
               Interfaces: ge-0/0/0.0
v1           224.1.1.9, *      1266
               Interfaces: ge-0/0/0.0
v1           224.1.1.11, *     1266
               Interfaces: ge-0/0/0.0
```

show igmp-snooping statistics

Syntax	show igmp-snooping statistics
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display IGMP snooping statistics.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Monitoring IGMP Snooping on page 72 • Configuring IGMP Snooping on page 68 • show igmp-snooping route on page 453 • show igmp-snooping vlans on page 457
List of Sample Output	show igmp-snooping statistics on page 456
Output Fields	Table 18 on page 455 lists the output fields for the show igmp-snooping statistics command. Output fields are listed in the approximate order in which they appear.

Table 18: show igmp-snooping statistics Output Fields

Field Name	Field Description
Bad length	IGMP packet has illegal or bad length.
Bad checksum	IGMP or IP checksum is incorrect.
Invalid interface	Packet was received through an invalid interface.
Not local	Number of packets received from senders that are not local.
Receive unknown	Unknown IGMP type.
Timed out	Number of timeouts for all multicast groups.
IGMP Type	Type of IGMP message (Queries , Reports , Leaves , or Other).
Received	Number of IGMP packets received.
Transmitted	Number of IGMP packets transmitted.
Recv Errors	Number of general receive errors.

Sample Output

show igmp-snooping statistics

```
user@switch> show igmp-snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 58
```

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

show igmp-snooping vlans

Syntax	<code>show igmp-snooping vlans</code> <code><brief detail></code> <code><vlan <i>vlan-id</i> <i>vlan-name</i>></code>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display IGMP snooping VLAN information.
Options	<p>none—Display general parameters.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>vlan <i>vlan-id</i> vlan <i>vlan-number</i>—(Optional) Display VLAN information for the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Monitoring IGMP Snooping on page 72 • Configuring IGMP Snooping on page 68 • show igmp-snooping route on page 453 • show igmp-snooping statistics on page 455
List of Sample Output	<p>show igmp-snooping vlans on page 458</p> <p>show igmp-snooping vlans vlan on page 458</p> <p>show igmp-snooping vlans vlan detail on page 458</p>
Output Fields	Table 19 on page 457 lists the output fields for the show igmp-snooping vlans command. Output fields are listed in the approximate order in which they appear.

Table 19: show igmp-snooping vlans Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All levels
IGMP-L2-Querier	Source address for IGMP snooping queries (if switch is an IGMP querier)	All levels
Interfaces	Number of interfaces in the VLAN.	All levels
Groups	Number of groups in the VLAN.	All levels
MRouters	Number of multicast routers associated with the VLAN.	All levels
Receivers	Number of host receivers in the VLAN.	All levels

Table 19: show igmp-snooping vlans Output Fields (*continued*)

Field Name	Field Description	Level of Output
Tag	Numerical identifier of the VLAN.	detail
tagged untagged	Interface participates in a tagged (802.1Q) or untagged (native) VLAN.	detail
vlan-interface	Internal VLAN interface identifier.	detail
Membership timeout	Membership timeout value.	detail
Querier timeout	Timeout value for interfaces dynamically marked as router or switch interfaces (interfaces that receive queries). When the querier timeout is reached, the switch marks the interface as a host interface.	detail
Interface	Name of the interface.	detail
Reporters	Number of dynamic groups on an interface.	detail

Sample Output

show igmp-snooping vlans

```

user@switch> show igmp-snooping vlans
VLAN      Interfaces Groups MRouters Receivers
default   0          0      0        0
v1         11         50      0        0
v10        1          0      0        0
v11        1          0      0        0
v180       3          0      1        0
v181       3          0      0        0
v182       3          0      0        0

```

show igmp-snooping vlans vlan

```

user@switch> show igmp-snooping vlans vlan v10
user@switch> show igmp-snooping vlans vlan v10
VLAN      Interfaces Groups MRouters Receivers
v10       1          0      0        0

```

show igmp-snooping vlans vlan detail

```

user@switch> show igmp-snooping vlans vlan v10 detail
VLAN: v10, Tag: 10, vlan-interface: vlan.10
      Interface: ge-0/0/10.0, tagged, Groups: 0
IGMP-L2-Querier: Stopped, SourceAddress: 10.10.1.2

```

CHAPTER 25

Operational Commands (PIM)

- 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 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 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>
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 486](#)

List of Sample Output [clear multicast bandwidth-admission on page 462](#)

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 463 Syntax (EX Series Switch and the QFX Series) on page 463
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	<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 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>
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 507
List of Sample Output	clear multicast scope on page 463
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 464 Syntax (EX Series Switch and the QFX Series) on page 464
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.
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 509
List of Sample Output	clear multicast sessions on page 464
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 465 Syntax (EX Series Switch and the QFX Series) on page 465
Syntax	clear multicast statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	clear multicast statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear IP multicast statistics.
Options	none —Clear multicast statistics for all supported address families on all interfaces. inet —(Optional) Clear multicast statistics for IPv4 family addresses. inet6 —(Optional) Clear multicast statistics for IPv6 family addresses. instance <i>instance-name</i> —(Optional) Clear multicast statistics for the specified instance. interface <i>interface-name</i> —(Optional) Clear multicast statistics on a specific interface. 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 multicast statistics
List of Sample Output	clear multicast statistics on page 465
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 466 Syntax (EX Series Switch and the QFX Series) on page 466
Syntax	<pre>clear pim join <group-address> <inet inet6> <instance instance-name> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>clear pim join <group-address> <inet inet6> <instance instance-name></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear the Protocol Independent Multicast (PIM) join and prune states.
Options	<p>none—Clear the PIM join and prune states for all groups, family addresses, and instances.</p> <p>group-address—(Optional) Clear the PIM join and prune states for a group address.</p> <p>inet inet6—(Optional) Clear the PIM join and prune states for IPv4 or IPv6 family addresses, respectively.</p> <p>instance instance-name—(Optional) Clear the join and prune states 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>
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	<ul style="list-style-type: none">• show pim join on page 520
List of Sample Output	clear pim join on page 467
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
```

clear pim register

List of Syntax	Syntax on page 468 Syntax (EX Series Switch and the QFX Series) on page 468 Syntax (PTX Series) on page 468
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	Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Protocol Independent Multicast (PIM) register message counters.
Options	<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 555](#)

List of Sample Output [clear pim register on page 469](#)

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 470 Syntax (EX Series Switch and the QFX Series) on page 470
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	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Protocol Independent Multicast (PIM) statistics.
Options	<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 555
List of Sample Output	clear pim statistics on page 471
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> <logical-system <i>logical-system-name</i> > <routing-instance <i>routing-instance-name</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 9.5 for SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 12.3 for the PTX Series.
Description	Display trace information about an IP multicast path.
Options	<i>source</i> —Source hostname or address. <i>logical-system (logical-system-name)</i> —(Optional) Perform this operation on a logical system. <i>routing-instance routing-instance-name</i> —(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 475
Output Fields	Table 20 on page 473 describes the output fields for the mtrace command. Output fields are listed in the approximate order in which they appear.

Table 20: 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.
<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.

Table 20: mtrace Output Fields (*continued*)

Field Name	Field Description
<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>`
 `<ttl ttl>`
 `<wait-time wait-time>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.3 for the QFX Series.

Description Display trace information about an IP multicast path from a source to this router or switch. If you specify a group address with this command, Junos OS returns additional information, such as packet rates and losses.

Options **brief | detail**—(Optional) Display the specified level of output.

extra-hops *extra-hops*—(Optional) Number of hops to take after reaching a nonresponsive router. You can specify a number between **0** and **255**.

group *group*—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

interval *interval*—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10** seconds.

loop—(Optional) Loop indefinitely, displaying rate and loss statistics.

max-hops *max-hops*—(Optional) Maximum hops to trace toward the source. The range of values is **0** through **255**. The default value is **32** hops.

max-queries *max-queries*—(Optional) Maximum number of query attempts for any hop. The range of values is 1 through **32**. The default is **3**.

multicast-response—(Optional) Always request the response using multicast.

no-resolve—(Optional) Do not attempt to display addresses symbolically.

no-router-alert—(Optional) Do not use the router-alert IP option.

response *response*—(Optional) Send trace response to a host or multicast address.

routing-instance *routing-instance-name*—(Optional) Trace a particular routing instance.

source *source*—Source hostname or address.

ttl *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 478](#)

Output Fields [Table 21 on page 477](#) describes the output fields for the **mtrace from-source** command. Output fields are listed in the approximate order in which they appear.

Table 21: 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 21: mtrace from-source Output Fields (*continued*)

Field Name	Field Description
Packet Statistics for Traffic From	Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop.
Receiver	IP address receiving the multicast.
Query source	IP address sending the mtrace query.

Sample Output

mtrace from-source

```

user@host> mtrace from-source source 192.1.4.2 group 225.1.1.1
Mtrace from 192.1.4.2 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
  0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-2 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
-3 hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.

Waiting to accumulate statistics...Results after 10 seconds:

Source      Response Dest    Overall    Packet Statistics For Traffic From
192.1.4.2   192.1.1.2  Packet    192.1.4.2 To 225.1.1.1
      v      ___/ rtt    2 ms      Rate      Lost/Sent = Pct  Rate
192.1.2.1
192.1.3.2   routerC.lab.mycompany.net
      v      ^      ttl    2          0/0    = --    0 pps
192.1.4.1
192.1.2.2   routerB.lab.mycompany.net
      v      \__  ttl    3          ?/0          0 pps
192.1.1.2   192.1.1.2
Receiver    Query Source

```

mtrace monitor

Syntax	mtrace monitor
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Listen passively for IP multicast responses. To exit the mtrace monitor command, type Ctrl+c.
Options	none —Trace the master instance.
Required Privilege Level	view
List of Sample Output	mtrace monitor on page 480
Output Fields	Table 22 on page 479 describes the output fields for the mtrace monitor command. Output fields are listed in the approximate order in which they appear.

Table 22: 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	<pre> 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> </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>
Description	Display trace information about a multicast path from this router or switch to a gateway router or switch.
Options	<p>gateway gateway—Send the trace query to a gateway multicast address.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>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.</p> <p>group group—(Optional) Group address for which to trace the path. The default group address is 0.0.0.0.</p> <p>interface interface-name—(Optional) Source address for sending the trace query.</p> <p>interval interval—(Optional) Number of seconds to wait before gathering statistics again. The default value is 10.</p> <p>loop—(Optional) Loop indefinitely, displaying rate and loss statistics.</p> <p>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.</p> <p>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.</p> <p>multicast-response—(Optional) Always request the response using multicast.</p> <p>no-resolve—(Optional) Do not attempt to display addresses symbolically.</p>

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 482](#)

Output Fields [Table 23 on page 482](#) describes the output fields for the **mtrace to-gateway** command. Output fields are listed in the approximate order in which they appear.

Table 23: 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 multicast flow-map

List of Syntax	Syntax on page 484 Syntax (EX Series Switch and the QFX Series) on page 484
Syntax	<pre>show multicast flow-map <brief detail> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>show multicast flow-map <brief detail></pre>
Release Information	<p>Command introduced in Junos OS Release 8.2.</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>
Description	Display configuration information about IP multicast flow maps.
Options	<p>none—Display configuration information about IP multicast flow maps on all systems.</p> <p>brief detail—(Optional) Display the specified level of output.</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 flow-map on page 485 show multicast flow-map detail on page 485
Output Fields	<p>Table 24 on page 484 describes the output fields for the show multicast flow-map command. Output fields are listed in the approximate order in which they appear.</p>

Table 24: 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 24: 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 486 Syntax (EX Series Switch and the QFX Series) on page 486
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>
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 487
Output Fields	<p>Table 25 on page 486 describes the output fields for the show multicast interface command. Output fields are listed in the approximate order in which they appear.</p>

Table 25: 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 25: 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.
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 489
Output Fields	Table 26 on page 488 describes the output fields for the show multicast minfo command. Output fields are listed in the approximate order in which they appear.

Table 26: 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 mriinfo

```
user@host> show multicast mriinfo 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 490 Syntax (EX Series Switch and the QFX Series) on page 490
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>
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 491 show multicast next-hops (Bidirectional PIM on page 491 show multicast next-hops brief on page 492 show multicast next-hops detail on page 492

Output Fields Table 27 on page 491 describes the output fields for the **show multicast next-hops** command. Output fields are listed in the approximate order in which they appear.

Table 27: 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 491](#).

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 493 Syntax (EX Series Switch and the QFX Series) on page 493
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 10.0.</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>
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
List of Sample Output	show multicast pim-to-igmp-proxy on page 494 show multicast pim-to-igmp-proxy instance on page 494
Output Fields	Table 28 on page 493 describes the output fields for the show multicast pim-to-igmp-proxy command. Output fields are listed in the order in which they appear.

Table 28: show multicast pim-to-igmp-proxy Output Fields

Field Name	Field Description
Instance	Routing instance. Default instance is master (inet.0 routing table).
Proxy state	State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: enabled or disabled .
<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 495 Syntax (EX Series Switch and the QFX Series) on page 495
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 496 show multicast pim-to-mld-proxy instance on page 496
Output Fields	Table 29 on page 495 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 29: 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 497](#)
 [Syntax \(EX Series Switch and the QFX Series\) on page 497](#)

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 Bidirectional PIM](#)

List of Sample Output [show multicast route on page 499](#)
[show multicast route \(Bidirectional PIM\) on page 500](#)
[show multicast route brief on page 500](#)
[show multicast route detail on page 501](#)
[show multicast route extensive \(Bidirectional PIM\) on page 501](#)
[show multicast route instance <instance-name> extensive on page 502](#)
[show multicast route summary on page 502](#)

Output Fields [Table 30 on page 498](#) describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

Table 30: 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
Session description	Name of the multicast session.	detail extensive

Table 30: show multicast route Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 499](#) or [show multicast route \(Bidirectional PIM\) on page 500](#).

show multicast route detail

```

user@host> show multicast route detail
Family: INET

Group: 228.0.0.0
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Unknown
  Statistics: 8 kbps, 100 pps, 45272 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Administratively Scoped
  Statistics: 0 kbps, 0 pps, 13404 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.70.15/32
  Upstream interface: so-1/0/0.0
  Downstream interface list:
    mt-1/1/0.1081344
  Session description: Administratively Scoped
  Statistics: 46 kbps, 1000 pps, 921077 packets

  Next-hop ID: 262143
  Upstream protocol: PIM

Family: INET6

```

show multicast route extensive (Bidirectional PIM)

```

user@host> show multicast route extensive
Family: INET

Group: 224.1.1.0/24
  Source: *
  Incoming interface list:
    lo0.0 ge-0/0/1.0
  Downstream interface list:
    ge-0/0/1.0
  Number of outgoing interfaces: 1
  Session description: NOB Cross media facilities
  Statistics: 0 kbps, 0 pps, 0 packets
  Next-hop ID: 2097153
  Incoming interface list ID: 585
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0

Group: 224.1.3.0/24

```

```

Source: *
Incoming interface list:
  lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
  ge-0/0/1.0
Number of outgoing interfaces: 1
Session description: NOB Cross media facilities
Statistics: 0 kbps, 0 pps, 0 packets
Next-hop ID: 2097153
Incoming interface list ID: 589
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0

```

Family: INET6

show multicast route instance <instance-name> extensive

```

user@host> show multicast route instance mvpn extensive
Family: INET

Group: 239.10.10.10
Source: 2.0.0.2/32
Upstream interface: xe-0/0/0.102
Downstream interface list:
  xe-10/3/0.0 xe-0/3/0.0 xe-0/0/0.106 xe-0/0/0.105
  xe-0/0/0.103 xe-0/0/0.104 xe-0/0/0.107 xe-0/0/0.108
Session description: Administratively Scoped
Statistics: 256 kbps, 3998 pps, 670150 packets
Next-hop ID: 1048579
Upstream protocol: MVPN
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 58
Uptime: 00:00:04

```

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 503 Syntax (EX Series Switch and the QFX Series) on page 503
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>
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 504 show multicast rpf inet6 on page 505 show multicast rpf prefix on page 506 show multicast rpf summary on page 506

Output Fields Table 31 on page 504 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

Table 31: 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 507 Syntax (EX Series Switch and the QFX Series) on page 507
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>
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 508 show multicast scope inet on page 508 show multicast scope inet6 on page 508
Output Fields	<p>Table 32 on page 507 describes the output fields for the show multicast scope command. Output fields are listed in the approximate order in which they appear.</p>

Table 32: show multicast scope Output Fields

Field Name	Field Description
Scope name	Name of the multicast scope.
Group Prefix	Range of multicast groups that are scoped.
Interface	Interface that is the boundary of the administrative scope.

Table 32: show multicast scope Output Fields (*continued*)

Field Name	Field Description
Resolve Rejects	Number of kernel resolve rejects.

Sample Output

show multicast scope

```
user@host> show multicast scope
```

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 509 Syntax (EX Series Switch and the QFX Series) on page 509
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.
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 510 show multicast sessions regular-expression detail on page 510
Output Fields	Table 33 on page 509 describes the output fields for the show multicast sessions command. Output fields are listed in the approximate order in which they appear.

Table 33: 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 512 Syntax (EX Series Switch and the QFX Series) on page 512
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.
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 513 show multicast usage brief on page 513 show multicast usage instance on page 513 show multicast usage detail on page 514
Output Fields	Table 34 on page 513 describes the output fields for the show multicast usage command. Output fields are listed in the approximate order in which they appear.

Table 34: 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 513](#).

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 515 Syntax (EX Series Switch and the QFX Series) on page 515
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>
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 516 show pim bootstrap instance on page 516
Output Fields	<p>Table 35 on page 515 describes the output fields for the show pim bootstrap command. Output fields are listed in the approximate order in which they appear.</p>

Table 35: show pim bootstrap Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
BSR	Bootstrap router.
Pri	Priority of the routing device as elected to be the bootstrap router.
Local address	Local routing device address.
Pri	Local routing device address priority to be elected as the bootstrap router.

Table 35: show pim bootstrap Output Fields (*continued*)

Field Name	Field Description
State	Local routing device election state: Candidate , Elected , or Ineligible .
Timeout	How long until the local routing device declares the bootstrap router to be unreachable, in seconds.

Sample Output

show pim bootstrap

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

BSR	Pri	Local address	Pri	State	Timeout
None	0	10.255.71.46	0	InEligible	0
feco:1:1:1:1:0:aff:785c	34	feco:1:1:1:1:0:aff:7c12	0	InEligible	0

show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A
```

BSR	Pri	Local address	Pri	State	Timeout
None	0	192.168.196.105	0	InEligible	0

show pim interfaces

List of Syntax	Syntax on page 517 Syntax (EX Series Switch and the QFX Series) on page 517
Syntax	<pre>show pim interfaces <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 pim interfaces <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>Support for bidirectional PIM added in Junos OS Release 12.1.</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 all routing instances.</p> <p>inet inet6—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about interfaces 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 interfaces on page 518
Output Fields	<p>Table 36 on page 517 describes the output fields for the show pim interfaces command. Output fields are listed in the approximate order in which they appear.</p>

Table 36: show pim interfaces Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Name	Interface name.
State	State of the interface. The state also is displayed in the show interfaces command.

Table 36: show pim interfaces Output Fields (*continued*)

Field Name	Field Description
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. (Not supported on QFX Series.) <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(sg)	Number of (s,g) join messages that have been seen on the interface.
JointCnt(*g)	Number of (*g) join messages that have been seen on the interface.
DR address	Address of the designated router.

Sample Output

show pim interfaces

```

user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,

```

Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name	Stat	Mode	IP	V	State	NbrCnt	JoinCnt(sg/*g)	DR address
ge-0/3/0.0	Up	S	4	2	NotDR,NotCap	1	0/0	40.0.0.3
ge-0/3/3.50	Up	S	4	2	DR,NotCap	1	9901/100	50.0.0.2
ge-0/3/3.51	Up	S	4	2	DR,NotCap	1	0/0	51.0.0.2
pe-1/2/0.32769	Up	S	4	2	P2P,NotCap	0	0/0	

show pim join

List of Syntax [Syntax on page 520](#)
 [Syntax \(EX Series Switch and the QFX Series\) on page 520](#)

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 466](#)
- [Example: Configuring Bidirectional PIM](#)

List of Sample Output

[show pim join summary on page 525](#)
[show pim join \(PIM Sparse Mode\) on page 525](#)
[show pim join \(Bidirectional PIM\) on page 526](#)
[show pim join inet6 on page 526](#)
[show pim join inet6 star-g on page 527](#)
[show pim join instance <instance-name> on page 527](#)
[show pim join detail on page 527](#)
[show pim join extensive \(PIM Sparse Mode\) on page 528](#)
[show pim join extensive \(Bidirectional PIM\) on page 529](#)
[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 530](#)
[show pim join instance <instance-name> extensive on page 530](#)
[show pim join extensive \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 531](#)
[show pim join extensive \(Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 532](#)
[show pim join summary on page 533](#)
[show pim join \(PIM Sparse Mode\) on page 534](#)
[show pim join \(Bidirectional PIM\) on page 534](#)
[show pim join inet6 on page 535](#)
[show pim join inet6 star-g on page 535](#)
[show pim join instance <instance-name> on page 535](#)
[show pim join detail on page 536](#)
[show pim join extensive \(PIM Sparse Mode\) on page 536](#)
[show pim join extensive \(Bidirectional PIM\) on page 537](#)

[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 538](#)

[show pim join instance <instance-name> extensive on page 538](#)

[show pim join extensive \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 539](#)

[show pim join extensive \(Multipoint LDP with Multicast-Only Fast Reroute\) on page 540](#)

Output Fields [Table 37 on page 522](#) describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

Table 37: 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
Flags	PIM flags: <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

Table 37: show pim join Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
Upstream state	<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
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

Table 37: show pim join Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
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
Assert Timeout		
Assert Timeout	Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null.	extensive

Table 37: show pim join Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 router 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: 224.1.1.1
  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 541 Syntax (EX Series Switch and the QFX Series) on page 541
Syntax	<pre>show pim neighbors <brief detail> <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 pim neighbors <brief detail> <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>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 543 show pim neighbors brief on page 543 show pim neighbors instance on page 543 show pim neighbors detail on page 543 show pim neighbors detail (With BFD) on page 544
Output Fields	<p>Table 38 on page 542 describes the output fields for the show pim neighbors command. Output fields are listed in the approximate order in which they appear.</p>

Table 38: 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 38: 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 543](#).

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 545 Syntax (EX Series Switch and the QFX Series) on page 545
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 548 show pim rps brief on page 548

[show pim rps <group-address> on page 548](#)
[show pim rps <group-address> \(Bidirectional PIM\) on page 548](#)
[show pim rps <group-address> \(PIM Dense Mode\) on page 549](#)
[show pim rps <group-address> \(SSM Range Without asm-override-ssm Configured\) on page 549](#)
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Sparse-Mode RP\) on page 549](#)
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Bidirectional RP\) on page 549](#)
[show pim rps instance on page 549](#)
[show pim rps extensive \(PIM Sparse Mode\) on page 549](#)
[show pim rps extensive \(Bidirectional PIM\) on page 550](#)
[show pim rps extensive \(PIM Anycast RP in Use\) on page 550](#)

Output Fields [Table 39 on page 546](#) describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

Table 39: 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 39: 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 39: 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	<p>If anycast RP is configured, the current register state for each group:</p> <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
10.10.1.3       static   bidir    150      None      2  224.1.3.0/24
                                   225.1.3.0/24
10.10.13.2      static   bidir    150      None      2  224.1.1.0/24
                                   225.1.1.0/24

```

show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see [show pim rps on page 548](#).

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 552 Syntax (EX Series Switch and the QFX Series) on page 552
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>
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 553 show pim source brief on page 553 show pim source detail on page 553
Output Fields	<p>Table 40 on page 553 describes the output fields for the show pim source command. Output fields are listed in the approximate order in which they appear.</p>

Table 40: 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 interface	RPF interface toward the source address.
Upstream Neighbor	Address of the RPF neighbor used to reach the source address.

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 553](#).

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 statistics

List of Syntax	Syntax on page 555 Syntax (EX Series Switch and the QFX Series) on page 555
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>
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 470
List of Sample Output	show pim statistics on page 561 show pim statistics inet interface <interface-name> on page 562 show pim statistics inet6 interface <interface-name> on page 563 show pim statistics interface <interface-name> on page 563
Output Fields	<p>Table 41 on page 556 describes the output fields for the show pim statistics command.</p> <p>Output fields are listed in the approximate order in which they appear.</p>

Table 41: 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.

Table 41: show pim statistics Output Fields (*continued*)

Field Name	Field Description
V2 State Refresh	PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh. State refresh is an extension to PIM-DM. It not supported in Junos OS.
V2 DF Election	PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election.
V1 Query	PIM version 1 query packets.
V1 Register	PIM version 1 register packets.
V1 Register Stop	PIM version 1 register stop packets.
V1 Join Prune	PIM version 1 join and prune packets.
V1 RP Reachability	PIM version 1 RP reachability packets.
V1 Assert	PIM version 1 assert packets.
V1 Graft	PIM version 1 graft packets.
V1 Graft Ack	PIM version 1 graft acknowledgment packets.
AutoRP Announce	Auto-RP announce packets.
AutoRP Mapping	Auto-RP mapping packets.
AutoRP Unknown type	Auto-RP packets with an unknown type.
Anycast Register	Auto-RP announce packets.
Anycast Register Stop	Auto-RP announce packets.
Global Statistics	Summary of PIM statistics for all interfaces.
Hello dropped on neighbor policy	Number of hello packets dropped because of a configured neighbor policy.
Unknown type	Number of PIM control packets received with an unknown type.
V1 Unknown type	Number of PIM version 1 control packets received with an unknown type.
Unknown Version	Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.

Table 41: show pim statistics Output Fields (*continued*)

Field Name	Field Description
Neighbor unknown	Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.
Bad Length	Number of PIM control packets received for which the packet size does not match the PIM length field in the packet.
Bad Checksum	Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet.
Bad Receive If	Number of PIM control packets received on an interface that does not have PIM configured.
Rx Bad Data	Number of PIM control packets received that contain data for TCP Bad register packets.
Rx Intf disabled	Number of PIM control packets received on an interface that has PIM disabled.
Rx V1 Require V2	Number of PIM version 1 control packets received on an interface configured for PIM version 2.
Rx V2 Require V1	Number of PIM version 2 control packets received on an interface configured for PIM version 1.
Rx Register not RP	Number of PIM register packets received when the router 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 router 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 router has no state.
Rx Join/Prune on upstream if	Number of join and prune messages received on the interface used to reach the upstream router, toward the RP.
Rx Join/Prune for invalid group	Number of join or prune messages received for invalid multicast group addresses.

Table 41: show pim statistics Output Fields (*continued*)

Field Name	Field Description
Rx Join/Prune messages dropped	Number of join and prune messages received and dropped.
Rx sparse join for dense group	Number of PIM sparse mode join messages received for a group that is configured for dense mode.
Rx Graft/Graft Ack no state	Number of graft and graft acknowledgment messages received for which the router or switch has no state.
Rx Graft on upstream if	Number of graft messages received on the interface used to reach the upstream router, 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 router has no state for the data type.
Rx RP no state	Number of PIM control packets received for which the router 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 router does not have an encapsulation interface.
No route upstream	Number of PIM control packets received when the router does not have a unicast route to the the interface used to reach the upstream router, 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 router has an RP mismatch.

Table 41: show pim statistics Output Fields (*continued*)

Field Name	Field Description
RP mode mismatch	RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages.
RPF neighbor unknown	Number of PIM control packets received for which the router has an unknown RPF neighbor for the source.
Rx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Tx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Embedded-RP invalid addr	Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.
Embedded-RP limit exceed	Number of times the limit configured with the maximum-rps statement is exceeded. The maximum-rps statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.
Embedded-RP added	<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 router:</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.

Table 41: show pim statistics Output Fields (*continued*)

Field Name	Field Description
Rx Bidir Join/Prune on non-DF if	Error counter for join and prune messages received on non-designated forwarder interfaces.

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            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	5
Rx Join/Prune messages dropped	0
Rx sparse join for dense group	0
Rx Graft/Graft Ack no state	0
Rx Graft on upstream if	0
Rx CRP not BSR	0
Rx BSR when BSR	0
Rx BSR not RPF if	0
Rx unknown hello opt	0
Rx data no state	0
Rx RP no state	0
Rx aggregate	0
Rx malformed packet	0
Rx illegal TTL	0
Rx illegal destination address	0
No RP	0
No register encap if	0
No route upstream	0
Nexthop Unusable	0
RP mismatch	0
RP mode mismatch	0
RPF neighbor unknown	0
Rx Joins/Prunes filtered	0
Tx Joins/Prunes filtered	0
Embedded-RP invalid addr	0
Embedded-RP limit exceed	0
Embedded-RP added	0
Embedded-RP removed	0
Rx Register msgs filtering drop	0
Tx Register msgs filtering drop	0
Rx Bidir Join/Prune on non-Bidir if	0
Rx Bidir Join/Prune on non-DF if	0

Sample Output

show pim statistics inet interface <interface-name>

```
user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET
```

PIM Interface statistics for ge-0/3/0.0

PIM Message type	Received	Sent	Rx errors
V2 Hello	0	4	0
V2 Register	0	0	0
V2 Register Stop	0	0	0
V2 Join Prune	0	0	0
V2 Bootstrap	0	0	0
V2 Assert	0	0	0
V2 Graft	0	0	0
V2 Graft Ack	0	0	0
V2 Candidate RP	0	0	0
V1 Query	0	0	0
V1 Register	0	0	0
V1 Register Stop	0	0	0
V1 Join Prune	0	0	0
V1 RP Reachability	0	0	0
V1 Assert	0	0	0
V1 Graft	0	0	0
V1 Graft Ack	0	0	0

AutoRP Announce	0	0	0
AutoRP Mapping	0	0	0
AutoRP Unknown type	0		
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

Sample Output

show pim statistics inet6 interface <interface-name>

```
user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6
```

PIM Interface statistics for ge-0/3/0.0

PIM Message type	Received	Sent	Rx errors
V2 Hello	0	4	0
V2 Register	0	0	0
V2 Register Stop	0	0	0
V2 Join Prune	0	0	0
V2 Bootstrap	0	0	0
V2 Assert	0	0	0
V2 Graft	0	0	0
V2 Graft Ack	0	0	0
V2 Candidate RP	0	0	0
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

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

CHAPTER 26

Operational Commands (MSDP)

- `clear msdp cache`
- `clear msdp statistics`
- `show msdp`
- `show msdp source`
- `show msdp source-active`
- `show msdp statistics`
- `test msdp`

clear msdp cache

Syntax	<code>clear msdp cache</code> <code><all></code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><peer <i>peer-address</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Clear the entries in the Multicast Source Discovery Protocol (MSDP) source-active cache.
Options	all — Clear all MSDP source-active cache entries in the master instance.. instance <i>instance-name</i> —(Optional) Clear entries for a specific MSDP instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. peer <i>peer-address</i> —(Optional) Clear the MSDP source-active cache entries learned from a specific peer.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show msdp source-active on page 572
List of Sample Output	clear msdp cache all on page 566
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear msdp cache all

```
user@host> clear msdp cache all
```

clear msdp statistics

Syntax	clear msdp statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <peer <i>peer-address</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear Multicast Source Discovery Protocol (MSDP) peer statistics.
Options	<p>none—Clear MSDP statistics for all peers.</p> <p>instance <i>instance-name</i>—(Optional) Clear statistics for the specified 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>peer <i>peer-address</i>—(Optional) Clear the statistics for the specified peer.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show msdp statistics on page 575
List of Sample Output	clear msdp statistics on page 567
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear msdp statistics

```
user@host> clear msdp statistics
```

show msdp

Syntax	<pre>show msdp <brief detail> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <peer <i>peer-address</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display Multicast Source Discovery Protocol (MSDP) information.
Options	<p>none—Display standard MSDP information for all routing instances.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified instance 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> <p>peer <i>peer-address</i>—(Optional) Display information about the specified peer only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show msdp source on page 570 • show msdp source-active on page 572 • show msdp statistics on page 575
List of Sample Output	<p>show msdp on page 569</p> <p>show msdp brief on page 569</p> <p>show msdp detail on page 569</p>
Output Fields	Table 42 on page 568 describes the output fields for the show msdp command. Output fields are listed in the approximate order in which they appear.

Table 42: show msdp Output Fields

Field Name	Field Description	Level of Output
Peer address	IP address of the peer.	All levels
Local address	Local address of the peer.	All levels
State	Status of the MSDP connection: Listen , Established , or Inactive .	All levels
Last up/down	Time at which the most recent peer-state change occurred.	All levels

Table 42: show msdp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Peer-Group	Peer group name.	All levels
SA Count	Number of source-active cache entries advertised by each peer that were accepted, compared to the number that were received, in the format <i>number-accepted/number-received</i> .	All levels
Peer Connect Retries	Number of peer connection retries.	detail
State timer expires	Number of seconds before another message is sent to a peer.	detail
Peer Times out	Number of seconds to wait for a response from the peer before the peer is declared unavailable.	detail
SA accepted	Number of entries in the source-active cache accepted from the peer.	detail
SA received	Number of entries in the source-active cache received by the peer.	detail

Sample Output

show msdp

```

user@host> show msdp
Peer address    Local address  State      Last up/down Peer-Group SA Count
198.32.8.193    198.32.8.195  Established 5d 19:25:44 North23 120/150
198.32.8.194    198.32.8.195  Established 3d 19:27:27 North23 300/345
198.32.8.196    198.32.8.195  Established 5d 19:39:36 North23 10/13
198.32.8.197    198.32.8.195  Established 5d 19:32:27 North23 5/6
198.32.8.198    198.32.8.195  Established 3d 19:33:04 North23 2305/3000

```

show msdp brief

The output for the **show msdp brief** command is identical to that for the **show msdp** command. For sample output, see [show msdp on page 569](#).

show msdp detail

```

user@host> show msdp detail
Peer: 10.255.70.15
Local address: 10.255.70.19
State: Established
Peer Connect Retries: 0
State timer expires: 22
Peer Times out: 49
SA accepted: 0
SA received: 0

```

show msdp source

Syntax	<code>show msdp source</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><source-address></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display multicast sources learned from Multicast Source Discovery Protocol (MSDP).
Options	none —Display standard MSDP source information for all routing instances. instance <i>instance-name</i> —(Optional) Display information for the specified instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. source-address —(Optional) IP address and optional prefix length. Display information for the specified source address only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show msdp on page 568• show msdp source-active on page 572• show msdp statistics on page 575
List of Sample Output	show msdp source on page 571

Output Fields Table 43 on page 571 describes the output fields for the **show msdp source** command. Output fields are listed in the approximate order in which they appear.

Table 43: show msdp source Output Fields

Field Name	Field Description
Source address	IP address of the source.
/Len	Length of the prefix for this IP address.
Type	Discovery method for this multicast source: <ul style="list-style-type: none"> • Configured—Source-active limit explicitly configured for this source. • Dynamic—Source-active limit established when this source was discovered.
Maximum	Source-active limit applied to this source.
Threshold	Source-active threshold applied to this source.
Exceeded	Number of source-active messages received from this source exceeding the established maximum.

Sample Output

show msdp source

```

user@host> show msdp source
Source address /Len  Type      Maximum  Threshold  Exceeded
0.0.0.0       /0    Configured    5         none        0
10.1.0.0      /16   Configured    500       none        0
10.1.1.1      /32   Configured    10000     none        0
10.1.1.2      /32   Dynamic       6936     none        0
10.1.5.5      /32   Dynamic       500       none        123
10.2.1.1      /32   Dynamic        2         none        0

```

show msdp source-active

Syntax	<code>show msdp source-active</code> <code><brief detail></code> <code><group <i>group</i>></code> <code><instance <i>instance-name</i>></code> <code><local></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><originator <i>originator</i>></code> <code><peer <i>peer-address</i>></code> <code><source <i>source-address</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the Multicast Source Discovery Protocol (MSDP) source-active cache.
Options	none —Display standard MSDP source-active cache information for all routing instances. brief detail —(Optional) Display the specified level of output. group <i>group</i> —(Optional) Display source-active cache information for the specified group. instance <i>instance-name</i> —(Optional) Display information for the specified instance. local —(Optional) Display all source-active caches originated by this router. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. originator <i>originator</i> —(Optional) Display information about the peer that originated the source-active cache entries. peer <i>peer-address</i> —(Optional) Display the source-active cache of the specified peer. source <i>source-address</i> —(Optional) Display the source-active cache of the specified source.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show msdp on page 568• show msdp source on page 570• show msdp statistics on page 575
List of Sample Output	show msdp source-active on page 573 show msdp source-active brief on page 574 show msdp source-active detail on page 574 show msdp source-active source on page 574

Output Fields Table 44 on page 573 describes the output fields for the **show msdp source-active** command. Output fields are listed in the approximate order in which they appear.

Table 44: show msdp source-active Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.
Group address	Multicast address of the group.
Source address	IP address of the source.
Peer address	IP address of the peer.
Originator	Router ID configured on the source of the rendezvous point (RP) that originated the message, or the loopback address when the router ID is not configured.
Flags	Flags: Accept , Reject , or Filtered .

Sample Output

show msdp source-active

```

user@host> show msdp source-active
Group address  Source address  Peer address  Originator  Flags
230.0.0.0     192.168.195.46  local        10.255.14.30  Accept
230.0.0.1     192.168.195.46  local        10.255.14.30  Accept
230.0.0.2     192.168.195.46  local        10.255.14.30  Accept
230.0.0.3     192.168.195.46  local        10.255.14.30  Accept
230.0.0.4     192.168.195.46  local        10.255.14.30  Accept

```

show msdp source-active brief

The output for the **show msdp source-active brief** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 573](#).

show msdp source-active detail

The output for the **show msdp source-active detail** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 573](#).

show msdp source-active source

```
user@host> show msdp source-active source 192.168.215.246
```

```
Global active source limit exceeded: 0
```

```
Global active source limit maximum: 25000
```

```
Global active source limit threshold: 24000
```

```
Global active source limit log-warning: 100
```

```
Global active source limit log interval: 0
```

Group address	Source address	Peer address	Originator	Flags
226.2.2.1	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.3	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.4	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.5	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.7	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.10	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.11	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.13	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.14	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.15	192.168.215.246	10.255.182.140	10.255.182.140	Accept

show msdp statistics

Syntax	show msdp statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <peer <i>peer-address</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display statistics about Multicast Source Discovery Protocol (MSDP) peers.
Options	none —Display statistics about all MSDP peers for all routing instances. instance <i>instance-name</i> —(Optional) Display statistics about a specific MSDP instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. peer <i>peer-address</i> —(Optional) Display statistics about a particular MSDP peer.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear msdp statistics on page 567
List of Sample Output	show msdp statistics on page 577 show msdp statistics peer on page 577
Output Fields	Table 45 on page 575 describes the output fields for the show msdp statistics command. Output fields are listed in the approximate order in which they appear.

Table 45: show msdp statistics Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.

Table 45: show msdp statistics Output Fields (*continued*)

Field Name	Field Description
Peer	Address of peer.
Last State Change	How long ago the peer state changed.
Last message received from the peer	How long ago the last message was received from the peer.
RPF Failures	Number of reverse path forwarding (RPF) failures.
Remote Closes	Number of times the remote peer closed.
Peer Timeouts	Number of peer timeouts.
SA messages sent	Number of source-active messages sent.
SA messages received	Number of source-active messages received.
SA request messages sent	Number of source-active request messages sent.
SA request messages received	Number of source-active request messages received.
SA response messages sent	Number of source-active response messages sent.
SA response messages received	Number of source-active response messages received.
SA messages with zero Entry Count received	Entry Count is a field within SA message that defines how many source/group tuples are present in the SA message. The counter is incremented each time an SA with an Entry Count of zero is received.
Active source exceeded	Number of times this peer has exceeded configured source-active limits.
Active source Maximum	Configured number of active source messages accepted by this peer.
Active source threshold	Configured threshold on this peer for applying random early discard (RED) to drop some but not all MSDP active source messages.
Active source log-warning	Configured threshold on this peer at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Active source log-interval	Time (in seconds) between consecutive log messages on this peer.
Keepalive messages sent	Number of keepalive messages sent.

Table 45: show msdp statistics Output Fields (*continued*)

Field Name	Field Description
Keepalive messages received	Number of keepalive messages received.
Unknown messages received	Number of unknown messages received.
Error messages received	Number of error messages received.

Sample Output

show msdp statistics

```

user@host> show msdp statistics
Global active source limit exceeded: 0
Global active source limit maximum: 10
Global active source limit threshold: 8
Global active source limit log-warning: 60
Global active source limit log interval: 60

Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA messages with zero Entry Count received: 0
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 0
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval 120
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0

```

show msdp statistics peer

```

user@host> show msdp statistics peer 10.255.182.140
Peer: 10.255.182.140
  Last State Change: 8:19:23 (00:01:08)
  Last message received from peer: 8:20:05 (00:00:26)
  RPF Failures: 0
  Remote Closes: 0
  Peer Timeouts: 0
  SA messages sent: 17
  SA messages received: 16
  SA request messages sent: 0
  SA request messages received: 0

```

```
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 20
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval: 120
Keepalive messages sent: 0
Keepalive messages received: 0
Unknown messages received: 0
Error messages received: 0
```

test msdp

Syntax	test msdp (dependent-peers <i>prefix</i> rpf-peer <i>originator</i>) <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Find Multicast Source Discovery Protocol (MSDP) peers.
Options	<p>dependent-peers <i>prefix</i>—Find downstream dependent MSDP peers.</p> <p>rpf-peer <i>originator</i>—Find the MSDP reverse-path-forwarding (RPF) peer for the originator.</p> <p>instance <i>instance-name</i>—(Optional) Find MDSP peers for the specified 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	test msdp dependent-peers on page 579
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

test msdp dependent-peers

```
user@host> test msdp dependent-peers 10.0.0.1/24
```


PART 2

Index

- [Index on page 583](#)

Index

Symbols

#, comments in configuration statements.....	xx
(), in syntax descriptions.....	xx
< >, in syntax descriptions.....	xx
[], in configuration statements.....	xx
{ }, in configuration statements.....	xx
(pipe), in syntax descriptions.....	xx

A

accept-remote-source statement	
usage guidelines.....	236
accounting statement	
IGMP.....	252
IGMP interface.....	252
active-source-limit statement.....	402
usage guidelines.....	238
address statement	
anycast RPs.....	323
usage guidelines.....	186, 244
local RPs.....	324
static RPs.....	325
usage guidelines.....	183
addresses	
multicast.....	28
administrative scoping.....	28
algorithm statement	
BFD authentication.....	326
all statement	
IGMP snooping.....	278, 300
anycast RP.....	192
overview.....	185
anycast-pim statement.....	327
usage guidelines.....	186, 244
asm-override-ssm statement.....	253, 395
assert (tracing flag).....	389
assert timeout	
configuring.....	222
assert-timeout statement.....	328
usage guidelines.....	222
authentication configuration	
BFD.....	135

authentication statement	
BFD.....	331
BFD protocol.....	329
authentication-key statement	
MSDP.....	403
auto-RP	
overview.....	195
auto-rp statement.....	330
usage guidelines.....	195

B

BFD	
authentication configuration.....	135
protocol.....	134
BFD authentication	
algorithm statement.....	326
authentication statement.....	329
key-chain statement.....	355
loose-check statement.....	358
bfd-liveness-detection statement	
PIM.....	331, 337
minimum-interval.....	362
threshold.....	387
transmit-interval.....	388
usage guidelines.....	134
bootstrap (tracing flag).....	389
bootstrap messages.....	201
bootstrap routers	
overview.....	201
bootstrap routers, displaying.....	515
bootstrap statement.....	332
bootstrap-export statement.....	333
bootstrap-import statement.....	334
bootstrap-priority statement.....	335
braces, in configuration statements.....	xx
brackets	
angle, in syntax descriptions.....	xx
square, in configuration statements.....	xx
BSR	
policy, import.....	350

C

cache (tracing flag).....	389
CBT	
defined.....	33
issues.....	125
clear igmp membership command.....	424
clear igmp statistics command.....	427
clear igmp-snooping membership command.....	448

clear igmp-snooping statistics command.....	449
clear msdp cache command.....	566
clear msdp statistics command.....	567
clear multicast bandwidth-admission command.....	461
clear multicast scope command.....	463
clear multicast sessions command.....	464
clear multicast statistics command.....	465
clear pim join command.....	466
clear pim register command.....	468
clear pim statistics command.....	470
comments, in configuration statements.....	xx
conventions text and syntax.....	xix
Core Based Trees See CBT	
curly braces, in configuration statements.....	xx
customer support.....	xxi
contacting JTAC.....	xxi
D	
data-encapsulation statement.....	404
usage guidelines.....	238
default-peer statement.....	405
usage guidelines.....	238
dense-groups statement.....	336
usage guidelines.....	161
designated router.....	146
detection-time statement PIM.....	337
disable statement IGMP.....	253
usage guidelines.....	63
IGMP snooping.....	279, 301
MLD usage guidelines.....	100
MSDP.....	406
PIM family.....	338
PIM interfaces.....	338
PIM protocol.....	338
distribution trees RPT.....	213
shared.....	213
documentation comments on.....	xxi
dr-election-on-p2p statement.....	339
PIM usage guidelines.....	133
dr-register-policy statement.....	339
usage guidelines.....	211

DVMRP defined.....	32
groups, displaying.....	512
dynamic IGMP statements promiscuous-mode interface.....	264

E	
embedded-rp statement.....	340
enable IGMP static group membership.....	52
enable MLD static group membership.....	89
event recording IGMP.....	59
MLD.....	96
exclude statement IGMP.....	254
usage guidelines.....	52
MLD usage guidelines.....	89
export statement MSDP.....	407
PIM.....	341, 342
configuring.....	208
PIM RP usage guidelines.....	201

F	
family statement bootstrap.....	343
local RP.....	345
PIM interfaces.....	344
PIM protocol.....	344
font conventions.....	xix
forwarding table multicast information, displaying.....	497
frames multicast snooping.....	29

G	
graceful restart PIM sparse-dense mode.....	161
graft (tracing flag) PIM.....	389
group joins limiting.....	60, 98
group membership SSM maps.....	175

group statement	
IGMP.....	255
usage guidelines.....	52
IGMP snooping.....	279, 301
MLD	
usage guidelines.....	89
MSDP.....	408
PIM RPF selection.....	346
group-count statement	
IGMP.....	256
usage guidelines.....	52
MLD	
usage guidelines.....	89
group-increment statement	
IGMP.....	256
usage guidelines.....	52
MLD	
usage guidelines.....	89
group-limit statement	
configuring.....	60
IGMP interface.....	257
IGMP snooping.....	280, 302
MLD	
usage guidelines.....	98
group-policy statement	
IGMP.....	257
usage guidelines.....	48
MLD	
usage guidelines.....	86
group-ranges statement.....	347
usage guidelines.....	183
groups	
DVMRP, displaying.....	512
IGMP membership, displaying.....	429
PIM	
general information, displaying.....	520
usage information, displaying.....	512
SSM.....	397
H	
hello (tracing flag)	
PIM.....	389
hello-interval statement	
PIM.....	348
usage guidelines.....	128
hold-time statement	
PIM.....	349
host-only-interface statement.....	281, 303

I	
IGMP.....	433
configuration statements.....	42
configuring.....	42
disabling.....	63
enabling.....	43, 258
event recording.....	59
group membership	
SSM maps for different groups to different	
sources.....	175
group membership, displaying.....	429
host-query message interval.....	45, 265
interface group limit.....	257
interfaces, displaying.....	435
last-member query interval.....	46, 266
overview.....	40
PIM-to-IGMP message translation information,	
displaying.....	493
query response interval.....	50, 267
robustness variable.....	51, 268
static group membership.....	52
statistics, displaying.....	439
tracing operations.....	61
version.....	44, 275
IGMP snooping	
configure the switch to be an IGMP	
querier.....	285, 307
enabling.....	282, 304
group limit.....	280, 302
group statement.....	279, 301
host-only interface.....	281, 303
host-query message interval.....	287, 310
last-member query interval.....	288, 311
query response interval.....	289, 312
source address.....	291, 314
static statement.....	279, 301
igmp statement.....	258
usage guidelines.....	43
IGMP statements	
promiscuous-mode	
interface.....	264
igmp-querier statement.....	281, 292, 303, 315
igmp-snooping statement.....	282, 304
IGMPv3.....	42
interoperability with older versions.....	42

immediate-leave statement	
IGMP.....	260
usage guidelines.....	47
IGMP snooping.....	283, 305
MLD	
usage guidelines.....	85
import statement	
bootstrap.....	350
usage guidelines.....	201
MSDP.....	409
PIM.....	351
usage guidelines.....	209
infinity statement.....	352
usage guidelines.....	224
interface lists.....	31
interface statement	
IGMP.....	261
usage guidelines.....	43
IGMP snooping.....	284, 285, 306, 307
MLD	
usage guidelines.....	82
PIM.....	353
usage guidelines.....	160
Internet Group Management Protocol See IGMP	
IP multicast	
announced sessions, displaying.....	509
bandwidth admission	
clearing.....	461
flow map information, displaying.....	484
forwarding table, displaying.....	497
interface information, displaying.....	486
network information, displaying.....	488
next-hop table, displaying.....	490
PIM-to-IGMP message translation information,	
displaying.....	493
PIM-to-MLD message translation information,	
displaying.....	495
RPF calculations, displaying.....	503
scope, clearing.....	463
scoped information, displaying.....	507
sessions, clearing.....	464
statistics	
clearing.....	465
tracing routes	
from the receiver to the source.....	473
from the source to the gateway	
router.....	481
from the source to the receiver.....	476
listen for responses.....	479
J	
join (tracing flag).....	389
join states, clearing PIM.....	466
join-load-balance statement.....	354
usage guidelines.....	147
join-prune-timeout statement.....	355
K	
keepalive (tracing flag)	
MSDP.....	419
key-chain statement	
BFD authentication.....	355
L	
I2-querier statement	
IGMP snooping.....	285, 307
leave (tracing flag)	
IGMP.....	273
load balancing	
for PIM join.....	147
local statement	
PIM.....	356
usage guidelines.....	181
local-address statement	
MSDP group.....	410
MSDP peer.....	410
PIM.....	357
loose-check statement	
BFD authentication.....	358
M	
manuals	
comments on.....	xxi
mapping-agent-election statement.....	359
usage guidelines.....	195
mappings	
SSM.....	399
maximum statement	
MSDP.....	411
usage guidelines.....	238
maximum-rps statement.....	360
maximum-transmit-rate statement	
IGMP.....	262
usage guidelines.....	52
MLD	
usage guidelines.....	88
mesh groups	
MSDP.....	238

minimum-interval	
PIM.....	362
minimum-interval statement	
PIM.....	361
usage guidelines.....	134
minimum-receive-interval statement	
PIM.....	331, 363
usage guidelines.....	134
MLD	
disabling.....	100
event recording.....	96
group membership	
SSM maps for different groups to different	
sources.....	175
host-query message interval.....	83
immediate-leave host removal	
configuring.....	85
last-member query interval.....	84
overview.....	75, 78
PIM-to-MLD message translation information,	
displaying.....	495
query response interval.....	84
robustness variable.....	87
static group membership.....	89
mld	
enabling.....	82
MLD snooping	
host-query message interval.....	287, 310
mld statement	
usage guidelines.....	81, 82
mld-snooping statement.....	308
mode statement	
MSDP.....	412
usage guidelines.....	238
PIM.....	363
usage guidelines.....	146, 161
MOSPF, defined.....	32
MSDP	
active source limit.....	402
maximum.....	411
per-source.....	417
threshold.....	418
authentication.....	403
cache entries, clearing.....	566
configuration statements.....	232
configuring.....	232
data-encapsulation.....	404
default peer.....	238, 405
enabling.....	413
general information, displaying.....	568
groups.....	408
local address.....	410
message source information, displaying.....	570
mode.....	412
peer statistics	
clearing.....	567
displaying.....	575
policy, routing.....	407, 409
routing tables.....	416
source-active cache, displaying.....	572
tracing operations.....	234
msdp statement.....	413
mt (tracing flag).....	389
mtrace (tracing flag)	
IGMP.....	61
mtrace command.....	473
mtrace from-source command.....	476
mtrace monitor command.....	479
mtrace to-gateway command.....	481
multicast	
addresses.....	28
administrative scoping.....	28
anycast RP.....	185
auto-RP.....	195
bootstrap router.....	201
defined.....	23
Layer 2 frames.....	29
leaf and branch.....	28
packet replication.....	35
protocols	
group membership.....	39
reverse-path forwarding (RPF).....	27
routing protocols.....	32
compared, table.....	34
shortest-path tree (SPT).....	27
snooping.....	29
SSM groups.....	397
SSM mapping.....	399
terminology.....	26
uses.....	25
multicast filters.....	205
MAC filters.....	206
MSDP SA messages.....	234
RP/DR register messages.....	206
configuring.....	211
multicast group joins	
limiting.....	60, 98
Multicast Listener Discovery See MLD	

Multicast Open Shortest Path First See	MOSPF
Multicast Source Discovery Protocol See	MSDP
multicast-router-interface statement	
IGMP snooping.....	286, 309
multiplier statement	
PIM.....	331, 364
usage guidelines.....	134

N

neighbor-policy statement.....	364
usage guidelines.....	207
next hops	
multicast entries, displaying.....	490
no-accounting statement	
IGMP.....	252
no-adaptation	
PIM.....	365
no-multicast-echo statement	
PIM	
usage guidelines.....	129
nsr-synchronization (tracing flag).....	390

O

olf-map statement	
IGMP.....	262
override-interval	
PIM.....	366
override-interval statement	
usage guidelines.....	151

P

packets (tracing flag)	
IGMP.....	274
PIM.....	390
parentheses, in syntax descriptions.....	xx
passive statement	
IGMP.....	263
peer statement	
MSDP.....	415
PIM	
anycast RP.....	327, 381
assert timeout.....	328, 384
configuring.....	222
background.....	125
BFD.....	134, 331, 361, 363, 364, 392
bidirectional mode	
defined.....	32
bootstrap messages import and export.....	201
bootstrap routers.....	201

bootstrap routers, displaying.....	515
configuring.....	128
dense mode.....	157, 160
defined.....	33
designated router.....	146
embedded RP.....	340
enabling.....	367
filters See	multicast filters
graceful restart	
sparse-dense mode.....	161
groups	
general information, displaying.....	520
usage information, displaying.....	512
hello interval.....	128
hold-time period.....	349
incoming join filter policy, applying.....	209
interfaces	
displaying.....	517
join load balancing	
configuring.....	147
join states, clearing.....	466
join suppression	
configuring.....	151
join-prune-timeout.....	355
maximum RPs.....	360
mixing modes.....	159
neighbors, displaying.....	541
network components.....	127
outgoing join filter policy, applying.....	208
overview.....	125
PIM-to-IGMP message translation information,	
displaying.....	493
PIM-to-MLD message translation information,	
displaying.....	495
policy, routing.....	351
prune states, clearing.....	466
register	
clearing.....	468
remote source.....	236
rendezvous point tree.....	214
routing tables.....	377
RPF, displaying source state.....	552
RPs.....	145, 181, 195, 213, 375, 378
anycast.....	327
anycast RP.....	185
displaying.....	545
embedded.....	340
mapping options.....	145
maximum.....	360

- source registration.....215
- SPT cutover control.....222
- sparse mode.....143, 146
 - defined.....33
- sparse-dense mode.....159, 336
 - defined.....33
- SSM.....164, 165, 172
- statistics
 - clearing.....470
 - displaying.....555
- version.....128, 146, 393
- pim statement.....367
 - usage guidelines.....128
- PIM-RP
 - SPT
 - configuring threshold cutover policy.....224
- policer, single-rate two-color
 - example.....175
- policy statement
 - SSM map.....396
- policy, import
 - BSR.....350
- policy, routing
 - MSDP.....407, 409
 - PIM.....351
 - PIM join filter.....208, 209
- prefix-list statement
 - PIM RPF selection.....370
- priority
 - PIM RPs.....373
- priority statement
 - bootstrap.....371
 - PIM.....372
 - usage guidelines.....132
 - usage guidelines.....201
- promiscuous-mode statement
 - IGMP
 - interface.....264
 - usage guidelines.....49
- propagation-delay statement.....374
 - usage guidelines.....151
- Protocol Independent Multicast *See* PIM
- protocols
 - group membership.....39
 - multicast routing.....32
 - compared, table.....34
- prune (tracing flag)
 - PIM.....390
- prune states, clearing PIM.....466

Q

- query-interval statement
 - IGMP.....265
 - usage guidelines.....45
 - IGMP snooping.....287, 310
 - MLD
 - usage guidelines.....83
 - MLD snooping.....287, 310
- query-last-member-interval statement
 - IGMP.....266
 - usage guidelines.....46
 - IGMP snooping.....288, 311
 - MLD
 - usage guidelines.....84
- query-response-interval statement
 - IGMP.....267
 - usage guidelines.....50
 - IGMP snooping.....289, 312
 - MLD
 - usage guidelines.....84

R

- real-time monitoring
 - IP multicast paths.....473
- register (tracing flag).....390
- register-probe-time statement.....375
- regular expressions
 - IP multicast scope
 - clearing.....463
 - IP multicast sessions
 - clearing.....464
 - displaying.....509
- rendezvous points *See* RPs *See* PIM and RP
- replication
 - multicast packet.....35
- report (tracing flag)
 - IGMP.....274
- reset-tracking-bit statement.....376
 - usage guidelines.....151
- reverse path forwarding *See* RPF
- reverse-path forwarding *See* RPF
- rib-group statement
 - MSDP.....416
 - PIM.....377
 - usage guidelines.....160

robust-count statement.....	290, 313	show msdp source command.....	570
IGMP.....	268	show msdp source-active command.....	572
usage guidelines.....	51	show msdp statistics command.....	575
MLD.....		show multicast flow-map command.....	484
usage guidelines.....	87	show multicast interface command.....	486
route (tracing flag)		show multicast minfo command.....	488
MSDP.....	419	show multicast next-hops command.....	490
routing solutions		show multicast pim-to-igmp-proxy	
multicast administrative scoping.....	28	command.....	493
multicast reverse-path forwarding (RPF).....	27	show multicast pim-to-mld-proxy command.....	495
multicast shortest-path tree (SPT).....	27	show multicast route command.....	497
routing tables		show multicast rpf command.....	503
MSDP.....	416	show multicast scope command.....	507
PIM.....	377	show multicast sessions command.....	509
RP		show multicast usage command.....	512
anycast.....	327	show pim bootstrap command.....	515
embedded.....	340	show pim interfaces command.....	517
rp (tracing flag).....	390	show pim join command.....	520
rp statement.....	378	show pim neighbors command.....	541
rp-register-policy statement.....	380	show pim rps command.....	545
usage guidelines.....	211	show pim source command.....	552
rp-set statement.....	381	show pim statistics command.....	555
usage guidelines.....	186, 244	show protocols igmp command.....	433
RPF		show system statistics igmp command.....	442
calculations, displaying.....	503	snooping	
PIM source state, displaying.....	552	IGMP and VLANs.....	69
RPF (reverse-path forwarding)		multicast.....	29
description.....	27	source filtering.....	42
rpf-selection statement		source statement	
PIM.....	382	IGMP.....	269
RPs		usage guidelines.....	52
displaying.....	545	MLD.....	
maximum.....	360	usage guidelines.....	89
RPT.....	213	MSDP.....	417
S		PIM RPF selection.....	365, 383
scoping, administrative.....	28	SSM.....	
shared trees.....	213	usage guidelines.....	169
shortest-path tree.....	27	source-active (tracing flag).....	419
shortest-path trees.....	218	source-active-request (tracing flag).....	419
See also SPT		source-active-response (tracing flag).....	419
show igmp group command.....	429	source-address statement	
show igmp interface command.....	435	IGMP snooping.....	291, 314
show igmp statistics command.....	439	source-count statement	
show igmp-snooping membership command.....	450	IGMP.....	270
show igmp-snooping route command.....	453	usage guidelines.....	52
show igmp-snooping statistics command.....	455	MLD.....	
show igmp-snooping vlans command.....	457	usage guidelines.....	89
show msdp command.....	568		

source-increment statement		
IGMP.....	271	
usage guidelines.....	52	
MLD		
usage guidelines.....	89	
source-specific multicast <i>See</i> SSM		
SPT.....	218	
configuring threshold cutover policy.....	224	
cutover control.....	222	
SPT (shortest-path tree).....	27	
spt-threshold statement.....	384	
usage guidelines.....	224	
SSM.....	164, 172	
configuring.....	167	
domains.....	169	
mapping.....	169	
SSM maps.....	175	
example.....	175	
SSM maps for different groups to different		
sources.....	175	
ssm-groups statement.....	397	
usage guidelines.....	172	
ssm-map statement		
IGMP.....	398	
usage guidelines.....	169	
MLD		
usage guidelines.....	169	
SSM.....	399	
usage guidelines.....	169	
ssm-map-policy statement		
IGMP interface.....	400	
static statement		
IGMP.....	272	
usage guidelines.....	52	
IGMP snooping.....	279, 293, 301, 316	
MLD		
usage guidelines.....	89	
PIM.....	385	
usage guidelines.....	183	
support, technical <i>See</i> technical support		
syntax conventions.....	xix	
T		
technical support		
contacting JTAC.....	xxi	
test msdp command.....	579	
threshold		
PIM.....	386, 387	
threshold statement		
MSDP.....	418	
usage guidelines.....	238	
traceoptions statement		
IGMP.....	273	
usage guidelines.....	61	
IGMP snooping.....	294, 317	
MSDP.....	419	
usage guidelines.....	234	
PIM.....	389	
usage guidelines.....	130	
tracing flags		
assert.....	389	
bootstrap.....	389	
cache, PIM.....	389	
graft		
PIM.....	389	
hello		
PIM.....	389	
join.....	389	
keepalive		
MSDP.....	419	
leave		
IGMP.....	273	
mt.....	389	
mtrace		
IGMP.....	61	
nsr-synchronization.....	390	
packets		
IGMP.....	274	
PIM.....	390	
prune		
PIM.....	390	
register.....	390	
report		
IGMP.....	274	
route		
MSDP.....	419	
rp.....	390	
source-active.....	419	
source-active-request.....	419	
source-active-response.....	419	
tracing IP multicast path		
from receiver to source.....	473	
from router to gateway.....	481	
from server to router.....	476	

tracing operations	
IGMP.....	61, 273
MSDP.....	234, 419
PIM.....	389
tracing routes	
from the receiver to the source.....	473
from the source to the gateway router.....	481
from the source to the receiver.....	476
monitoring.....	479
transmit-interval	
PIM.....	388

V

version statement	
BFD.....	392
IGMP.....	275
usage guidelines.....	44
MLD	
usage guidelines.....	83
PIM.....	393
usage guidelines.....	128, 134, 146, 183
vlan statement	
IGMP snooping.....	297, 320
usage guidelines.....	69
VLANs	
IGMP snooping.....	69

W

wildcard-source statement	
PIM RPF selection.....	394