JUNOS® Software

Multicast Protocols Configuration Guide

Release 10.0

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

Revision 1
Published: 2009-10-09

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

Juniper Networks, the Juniper Networks logo, JUNOS, NetScreen, ScreenOS, and Steel-Belted Radius are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOSe is a trademark 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.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

JUNOS® Software Multicast Protocols Configuration Guide
Release 10.0
Copyright © 2009, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Writing: Roy Spencer, Walter Goralski, Mark Barnard
Editing: Nancy Kurahashi, Dawn Spencer
Illustration: Faith Bradford
Cover Design: Edmonds Design

Revision History
October 2009—R1 JUNOS 10.0

The information in this document is current as of the date listed in the revision history.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. The JUNOS Software 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

READ THIS END USER LICENSE AGREEMENT ("AGREEMENT") BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer's principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer's principal office is located outside the Americas) (such applicable entity being referred to herein as "Juniper"), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. "Software" also includes updates, upgrades and new releases of such software. "Embedded Software" means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.

b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.

c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.

d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.

e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

5. **Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.

7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.

8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.

9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.

10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.

11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.

12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14(ALT III) as applicable.

13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.

14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at http://www.gnu.org/licenses/gpl.html, and a copy of the LGPL at http://www.gnu.org/licenses/lgpl.html.

15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous

agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentés confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattaché, soient redigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Abbreviated Table of Contents

		About This Guide	xxxi
Part 1		**Introduction to Multicast**	
	Chapter 1	Multicast Overview	3
	Chapter 2	IP Multicast Overview	17
Part 2		**IP Multicast Configuration**	
	Chapter 3	Introduction to PIM	27
	Chapter 4	Complete Multicast Configuration Statements	33
Part 3		**IGMP**	
	Chapter 5	IGMP Overview	53
	Chapter 6	IGMP Configuration Guidelines	55
	Chapter 7	Summary of IGMP Configuration Statements	75
Part 4		**MLD**	
	Chapter 8	MLD Overview Introduction	95
	Chapter 9	MLD Configuration Guidelines	99
	Chapter 10	Summary of MLD Configuration Statements	117
Part 5		**SAP and SDP**	
	Chapter 11	SAP Overview	139
	Chapter 12	SAP Configuration Guidelines Introduction	141
	Chapter 13	Summary of SAP Configuration Statements	143
Part 6		**PGM**	
	Chapter 14	PGM Overview	149
	Chapter 15	PGM Configuration Guidelines Overview	155
	Chapter 16	Summary of PGM Configuration Statements	157
Part 7		**Multicast Routing Instances**	
	Chapter 17	Multicast over Layer 3 Draft-Rosen VPNs Overview	163

Chapter 18 Draft-Rosen Multicast VPN Configuration Guidelines 169
Chapter 19 Summary of Draft-Rosen MVPN Configuration Statements 177

Part 8 Multicast Routing Options

Chapter 20 Multicast Administrative Scoping Configuration Guidelines 193
Chapter 21 Multicast Reverse Path Forwarding Configuration Guidelines 199
Chapter 22 Source-Specific Multicast Configuration Guidelines 203
Chapter 23 Flow Map Configuration Guidelines 211
Chapter 24 Bandwidth Management Configuration Guidelines 215
Chapter 25 Multicast Forwarding Cache Configuration Guidelines 229
Chapter 26 Ingress PE Redundancy Configuration Guidelines 235
Chapter 27 Configuring PIM-to-IGMP and PIM-to-MLD Message Translation 239
Chapter 28 Summary of Multicast Routing Options Configuration Statements 245

Part 9 Multicast Snooping

Chapter 29 Multicast Snooping Overview 275
Chapter 30 Multicast Snooping Configuration Guidelines 277
Chapter 31 Summary of Multicast Snooping Configuration Statements 283
Chapter 32 IGMP Snooping Overview 289
Chapter 33 IGMP Snooping Configuration Guidelines 293
Chapter 34 Summary of IGMP Snooping Configuration Statements 299

Part 10 DVMRP

Chapter 35 DVMRP Overview 317
Chapter 36 DVMRP Configuration Guidelines 319
Chapter 37 Summary of DVMRP Configuration Statements 329

Part 11 PIM

Chapter 38 PIM Overview 341
Chapter 39 PIM Mode-Independent Configuration Guidelines 369
Chapter 40 PIM Dense and Sparse-Dense Mode Configuration Guidelines 391
Chapter 41 PIM Sparse Mode Configuration Guidelines 393
Chapter 42 PIM Configuration Examples 421
Chapter 43 Summary of PIM Configuration Statements 445

Part 12 MSDP

Chapter 44 MSDP Introduction 495
Chapter 45 MSDP Configuration Guidelines 497
Chapter 46 Summary of MSDP Configuration Statements 515

Part 13 Index

Index 539
Index of Statements and Commands 551

Table of Contents

About This Guide xxxi

JUNOS Documentation and Release Notes xxxi
Objectives xxxii
Audience xxxii
Supported Platforms xxxii
Using the Indexes xxxiii
Using the Examples in This Manual xxxiii
    Merging a Full Example xxxiii
    Merging a Snippet xxxiv
Documentation Conventions xxxiv
Documentation Feedback xxxvi
Requesting Technical Support xxxvi
    Self-Help Online Tools and Resources xxxvii
    Opening a Case with JTAC xxxvii

Part 1 **Introduction to Multicast**

Chapter 1 **Multicast Overview** 3

Comparing Multicast to Unicast 3
IP Multicast Uses 5
IP Multicast Terminology 6
Multicast Leaf and Branch Terminology 7
Dense and Sparse Modes for Multicast Networks 7
IP Multicast Building Blocks Overview 8
IP Multicast Addressing 8
Multicast Administrative Scoping 9
Multicast Interface Lists 9
Multicast Reverse-Path Forwarding 10
Multicast Shortest-Path Tree 11
Multicast Rendezvous Points, Shared Trees, and the Rendezvous-Point Trees 12
Multicast RPF Checks 13
Populating the Multicast RPF Table 14
Protocols for Multicast 14
Multicast Group Membership Protocols 14
Multicast Routing Protocols 15

Chapter 2 **IP Multicast Overview** **17**

JUNOS Software Multicast Standards 17
IP Multicast Specifications 18
Multicast Overview 19
Multicast Addresses 20
Multicast Redundancy 20
T Series Router Multicast Performance 21
Layer 2 Frames and IPv4 Multicast Addresses 21
Multicast Snooping Overview 23

Part 2 **IP Multicast Configuration**

Chapter 3 **Introduction to PIM** **27**

PIM Background 27
Basic PIM Network Components 28
PIM Modes of Operation 28
PIM Dense Mode Overview 29
PIM Sparse Mode Overview 29
PIM Source-Specific Mode Overview 30
Mixing PIM Sparse and Dense Modes 31

Chapter 4 **Complete Multicast Configuration Statements** **33**

Multicast Configuration Statements Overview 33
[edit protocols] Hierarchy Level 33
[edit routing-instances] Hierarchy Level 38
[edit routing-options] Hierarchy Level 40
[edit logical-systems protocols] Hierarchy Level 41
[edit logical-systems routing-instances] Hierarchy Level 45
[edit logical-systems routing-options] Hierarchy Level 48
[edit bridge-domains bridge-domain-name protocols] Hierarchy Level 49

Part 3 **IGMP**

Chapter 5 **IGMP Overview** **53**

Chapter 6 **IGMP Configuration Guidelines** **55**

Configuring IGMP 55
Minimum IGMP Configuration 56
Enabling IGMP 57

Modifying the IGMP Host-Query Message Interval 57
Modifying the IGMP Query Response Interval 57
Specifying Immediate-Leave Host Removal for IGMP 58
Example: IGMP Immediate Leave 58
Filtering Unwanted IGMP Reports at the IGMP Interface Level 59
Example: IGMP Report Filtering 59
Accepting IGMP Messages from Remote Subnetworks 60
Example: IGMP Promiscuous Mode 60
Modifying the IGMP Last-Member Query Interval 60
Modifying the IGMP Robustness Variable 61
Limiting the Maximum IGMP Message Rate 62
Changing the IGMP Version 62
Enabling IGMP Static Group Membership 63
Recording IGMP Join and Leave Events 69
Enabling IGMP Accounting on the Entire Routing System 70
Enabling or Disabling IGMP Accounting on Individual Interfaces 70
Example: Recording and Archiving IGMP Join and Leave Events 71
Tracing IGMP Protocol Traffic 71
Example: Tracing IGMP Protocol Traffic 72
Disabling IGMP 73
IGMP and Nonstop Active Routing 73

Chapter 7 Summary of IGMP Configuration Statements 75

accounting (Per Interface) 75
accounting (Protocol) 75
disable 76
exclude 76
group 77
group-count 77
group-increment 78
group-policy 78
igmp 79
immediate-leave 80
interface 81
maximum-transmit-rate 82
no-accounting 82
oif-map 83
passive (IGMP) 84
promiscuous-mode 85
query-interval 85
query-last-member-interval 86
query-response-interval 86
robust-count 87
source 87
source-count 88
source-increment 88
ssm-map 89
static 89

traceoptions 90
version 92

Part 4 MLD

Chapter 8 MLD Overview Introduction 95

MLD Overview 95

Chapter 9 MLD Configuration Guidelines 99

Configuring MLD 99
Minimum MLD Configuration 100
Enabling MLD 100
Modifying the MLD Version 101
Modifying the MLD Host-Query Message Interval 101
Modifying the MLD Query Response Interval 101
Modifying the MLD Last-Member Query Interval 102
Specifying Immediate-Leave Host Removal for MLD 102
Example: MLD Immediate Leave 103
Filtering Unwanted MLD Reports at the MLD Interface Level 103
Example: MLD Report Filtering 104
Modifying the MLD Robustness Variable 104
Limiting the Maximum MLD Message Rate 105
Enabling MLD Static Group Membership 105
Recording MLD Join and Leave Events 112
Enabling MLD Accounting on the Entire Routing System 112
Enabling or Disabling MLD Accounting on Individual Interfaces 113
Example: Recording and Archiving MLD Join and Leave Events 113
Tracing MLD Protocol Traffic 114
Example: Tracing MLD Protocol Traffic 115
Disabling MLD 115

Chapter 10 Summary of MLD Configuration Statements 117

accounting (Per Interface) 117
accounting (Protocol) 117
disable 118
exclude 118
group 119
group-count 120
group-increment 120
group-policy 121
immediate-leave 122
interface 123
maximum-transmit-rate 124
mld 125

no-accounting 126
oif-map 127
passive (MLD) 127
query-interval 128
query-last-member-interval 128
query-response-interval 129
robust-count 129
source 130
source-count 130
source-increment 131
ssm-map 131
static 132
traceoptions 133
version 135

Part 5 **SAP and SDP**

Chapter 11 **SAP Overview** **139**

Chapter 12 **SAP Configuration Guidelines Introduction** **141**

SAP Configuration Guidelines 141

Chapter 13 **Summary of SAP Configuration Statements** **143**

disable 143
listen 144
sap 145

Part 6 **PGM**

Chapter 14 **PGM Overview** **149**

Pragmatic General Multicast Overview 149
PGM Architecture and PGM Routers 150
PGM-Enabled Source 151
PGM-Enabled Receivers 151
PGM-Enabled Routers 152

Chapter 15 **PGM Configuration Guidelines Overview** **155**

PGM Configuration Guidelines 155

Chapter 16 **Summary of PGM Configuration Statements** **157**

pgm 157
traceoptions 158

Part 7 **Multicast Routing Instances**

Chapter 17 **Multicast over Layer 3 Draft-Rosen VPNs Overview** **163**

Draft-Rosen Multicast VPNs Overview 163
Data MDTs Overview 164
Data MDT Characteristics 165
Draft-Rosen Multicast VPN Autodiscovery 166
Draft-Rosen Multicast VPN Control Plane 166
Interoperating with Other Vendors 167

Chapter 18 **Draft-Rosen Multicast VPN Configuration Guidelines** **169**

Configuring Draft-Rosen Multicast VPNs 169
Configuring Data MDTs and Provider Tunnels 169
- Configuring the Data MDT Group Range 170
- Configuring the Data MDT Threshold Parameters 171
- Limiting the Maximum Number of Data MDTs 172

Enabling Dynamic Reuse of Data MDT Group Addresses 173
Configuring the Provider Tunnel SSM Group Address 173
Configuring Draft-Rosen Multicast VPN Autodiscovery 173
Controlling the Single Forwarder Election 175
Configuring the Draft-Rosen Multicast VPN Control Plane 175
Configuring Multicast for Virtual Routers 176

Chapter 19 **Summary of Draft-Rosen MVPN Configuration Statements** **177**

autodiscovery 177
autodiscovery-only 178
data-mdt-reuse 179
group 179
group-address 180
group-range 180
inclusive 181
inet-mdt 181
intra-as 182
mdt 183
mvpn (Control Plane Autodiscovery) 184
mvpn (Routing Instance) 185
pim-ssm 185

rate 186
signaling 187
source 188
threshold 189
tunnel-limit 190
unicast-umh-election 190

Part 8 **Multicast Routing Options**

Chapter 20 **Multicast Administrative Scoping Configuration Guidelines 193**

Multicast Scoping Overview 193
Creating a Named Scope for Multicast Scoping 194
Using a Scope Policy for Multicast Scoping 195
Example: Creating a Named Scope for Multicast Scoping 196
Example: Using a Scope Policy for Multicast Scoping 197

Chapter 21 **Multicast Reverse Path Forwarding Configuration Guidelines 199**

Multicast RPF Configuration Guidelines 199
Configuring RPF Policies 200
Example: Configuring RPF Policies 200

Chapter 22 **Source-Specific Multicast Configuration Guidelines 203**

Source-Specific Multicast Groups Overview 203
Configuring Source-Specific Multicast Groups 204
Source-Specific Multicast Examples 206
Example: Configuring an SSM-Only Domain 206
Example: Configuring PIM SSM on a Network 206
Example: Configuring SSM Mapping 208

Chapter 23 **Flow Map Configuration Guidelines 211**

Creating a Multicast Flow Map 211
    Creating a Flow Map Policy 211
    Defining Flow Maps 212
Defining Flow Properties 212
    Defining Bandwidth for Multicast Flows 212
    Defining Forwarding Cache Timeout 212
    Specifying Redundant Flow Sources 212
Displaying Flow Maps 213

Chapter 24 **Bandwidth Management Configuration Guidelines** **215**

Bandwidth Management Overview 215
Bandwidth Management and PIM Graceful Restart 216
Bandwidth Management and Source Redundancy 216
Logical Systems and Bandwidth Oversubscription 217
How Interface Bandwidth Becomes Oversubscribed 217
How Interface Bandwidth Becomes Available Again 217
Readmitting or Removing Interfaces 217
Defining Interface Bandwidth Maximums 218
Defining Bandwidth for Multicast Flows 218
Examples: Defining Bandwidths 219
Example: Configuring Maximum Multicast Bandwidth on an Interface 219
Example: Configuring Bandwidth for Individual Flows 220
Configuring Multicast Outgoing Interface Mapping 220
Configuring an OIF Map 221
Associating OIF Maps with Interfaces 222
Disabling QoS Adjustment 223
Configuring Passive Mode for IGMP or MLD 224
Configuring Multicast Routing Over IP Demux Interfaces 225
Managing Subscriber Overcommitment 226
Configuring Multicast Egress Class of Service 227

Chapter 25 **Multicast Forwarding Cache Configuration Guidelines** **229**

Multicast Forwarding Cache Overview 229
Configuring General Multicast Forwarding Cache Properties 229
Configuring Multicast Forwarding Cache Properties for Flow Maps 230
Examples: Configuring Multicast Forwarding Cache Properties 231
Configuring Forwarding Cache Properties at the Multicast Level 231
Configuring Forwarding Cache Properties at the Flow Map Level 232
Displaying the Cache Timeout 233

Chapter 26 **Ingress PE Redundancy Configuration Guidelines** **235**

Ingress PE Redundancy Configuration Overview 235
Configuring Ingress PE Redundancy 235
Example: Ingress PE Redundancy 237

Chapter 27 **Configuring PIM-to-IGMP and PIM-to-MLD Message Translation** **239**

PIM-to-IGMP and PIM-to-MLD Message Translation Overview 239
Configuring PIM-to-IGMP Message Translation 241
Configuring PIM-to-MLD Message Translation 242

Chapter 28 **Summary of Multicast Routing Options Configuration Statements** **245**

asm-override-ssm 245
backup-pe-group 246
backups 247
bandwidth 248
flow-map 249
forwarding-cache (Flow Maps) 250
forwarding-cache (Multicast) 251
interface (Routing Options) 252
interface (Scoping) 253
local-address 253
maximum-bandwidth 254
multicast 255
no-qos-adjust 257
pim-to-igmp-proxy 258
pim-to-mld-proxy 259
policy (Flow Maps) 260
policy (SSM Maps) 261
prefix 261
redundant-sources 262
reverse-oif-mapping 263
rpf-check-policy 264
scope 264
scope-policy 265
source 266
ssm-groups 266
ssm-map 267
subscriber-leave-timer 268
threshold 269
timeout (Flow Maps) 270
timeout (Multicast) 271
upstream-interface 272

Part 9 **Multicast Snooping**

Chapter 29 **Multicast Snooping Overview** **275**

Chapter 30 **Multicast Snooping Configuration Guidelines** **277**

Configuring Multicast Snooping 277
Configuring Forwarding Cache Snooping Options 278
Configuring Flood Groups for Snooping 279
Configuring Graceful Restart for Snooping 279
Ignoring Spanning Tree Topology Change Messages 280

Chapter 31 **Summary of Multicast Snooping Configuration Statements** **283**

flood-groups 283
forwarding-cache 284
graceful-restart 285
ignore-stp-topology-change 285
multicast-snooping-options 286
threshold 287

Chapter 32 **IGMP Snooping Overview** **289**

Introduction to IGMP Snooping 289
IGMP Snooping Interfaces and Forwarding 290
IGMP Snooping and Proxies 291
Multicast-Router Interfaces and IGMP Snooping Proxy Mode 292
Host-Side Interfaces and IGMP Snooping Proxy Mode 292
IGMP Snooping and Bridge Domains 292

Chapter 33 **IGMP Snooping Configuration Guidelines** **293**

Configuring IGMP Snooping 293
Configuring IGMP Snooping Proxy Mode 294
Configuring IGMP Snooping Immediate Leave 294
Configuring General IGMP Snooping Parameters 295
Configuring IGMP Snooping Interfaces 297
Configuring VLAN-Specific IGMP Snooping Parameters 298
Tracing IGMP Snooping Operations 298

Chapter 34 **Summary of IGMP Snooping Configuration Statements** **299**

group 299
group-limit 300
host-only-interface 301
igmp-snooping 302
immediate-leave 303
interface 304
multicast-router-interface 305
proxy 306
query-interval 307
query-last-member-interval 308
query-response-interval 309
robust-count 310
source 310
source-address 311
static 312
vlan 313

Part 10 DVMRP

Chapter 35 DVMRP Overview 317

Chapter 36 DVMRP Configuration Guidelines 319

Configuring DVMRP 319
Minimum DVMRP Configuration 320
Creating Routing Tables for DVMRP Routes 321
Enabling DVMRP 321
Modifying the DVMRP Hold-Time Period 322
Modifying the DVMRP Metric Value 322
Disabling DVMRP on an Interface 323
Configuring DVMRP Routing Policy 323
Configuring DVMRP Routing Modes 324
Tracing DVMRP Protocol Traffic 324
DVMRP Configuration Examples 325
Example: Tracing DVMRP Protocol Traffic 325
Example: Configuring DVMRP 326
Example: Configuring DVMRP to Announce Unicast Routes 326

Chapter 37 Summary of DVMRP Configuration Statements 329

disable 329
dvmrp 330
export 331
hold-time 331
import 332
interface 332
metric 333
mode 333
rib-group 334
traceoptions 335

Part 11 PIM

Chapter 38 PIM Overview 341

Protocol Independent Multicast Overview 341
PIM Sparse Mode 342
Designated Router 344
Rendezvous Point 344
RP Mapping Options 344
  Static Configuration 345
  Anycast RP 345

Auto-RP 345
Bootstrap Router 346
Building an RPT Between RP and Receivers 346
PIM Sparse-Mode Source Registration 347
PIM Sparse-Mode SPT Cutover 350
SPT Cutover 350
SPT Cutover Control 354
PIM SSM 354
PIM Dense Mode 355
PIM Sparse-Dense Mode 357
RP Mapping with Anycast RP 357
Multicast over Layer 3 VPNs 358
Dual PIM Multicast VPNs 358
MBGP-Based Multicast VPNs 359
Tunnel Services PICs and Multicast 362
Filtering Multicast Messages 363
Filtering MAC Addresses 364
Filtering RP/DR Register Messages 364
Filtering MSDP SA Messages 365
Embedded RP for IPv6 Multicast 366

Chapter 39 **PIM Mode-Independent Configuration Guidelines** **369**

PIM Mode-Independent Configuration Overview 369
PIM Configuration Statements 370
Changing the PIM Version 372
Configuring Interface Priority to Become the PIM Designated Router 372
Configuring PIM Designated Router Election on Point-to-Point Links 373
Modifying the PIM Hello Interval 373
Configuring Interface-Level PIM Neighbor Policies 373
Configuring a PIM RPF Routing Table 374
Filtering Outgoing PIM Join Messages 376
Filtering Incoming PIM Join Messages 377
Preserving Multicast Performance by Disabling Response to the ping Utility 378
Configuring the BFD Protocol for PIM 378
Overview of BFD Authentication for PIM 379
BFD Authentication Algorithms 380
Security Authentication Keychains 380
Strict Versus Loose Authentication 381
Configuring BFD Authentication for PIM 381
Configuring BFD Authentication Parameters 382
Viewing Authentication Information for BFD Sessions 383
Configuring Nonstop Active Routing with PIM 384
Disabling PIM 385
Disabling the PIM Protocol 386
Disabling PIM on an Interface 386
Disabling PIM for a Family 387
Disabling PIM for a Rendezvous Point 388
Configuring PIM Trace Options 388

Chapter 40 **PIM Dense and Sparse-Dense Mode Configuration Guidelines** **391**

Configuring PIM Dense Mode Properties 391
Configuring PIM Sparse-Dense Mode Properties 392

Chapter 41 **PIM Sparse Mode Configuration Guidelines** **393**

Minimum PIM Sparse Mode Configuration 393
Enabling PIM Sparse Mode 394
Configuring PIM Sparse Mode Graceful Restart 395
Configuring Local PIM RPs 396
    Configuring the IP Protocol Family 397
    Configuring the Local RP Address 398
    Configuring the Routing Device's RP Priority 398
    Configuring the Groups for Which the Routing Device Is the RP 398
    Modifying the Local RP Hold Time 399
Configuring Static PIM RPs 399
Configuring PIM Bootstrap Properties 400
    Configuring the Routing Device's IPv4 Bootstrap Router Priority 400
    Filtering PIM IPv4 Bootstrap Messages 400
    Configuring the Routing Device's Bootstrap Router Priority 401
    Filtering PIM Bootstrap Messages 402
Configuring PIM Auto-RP 402
    Configuring Auto-RP Announcement, Mapping, and Discovery 403
    Configuring Auto-RP Mapping Agent Election 406
Configuring Register Message Filtering on a PIM RP or DR 407
Configuring PIM Join Load Balancing 408
Modifying the IGMP Join State Timeout 412
Configuring PIM Embedded RP for IPv6 412
Configuring the PIM Assert Timeout 413
Configuring the PIM SPT Threshold Policy 413
    SPT Threshold Policy Configuration Changes 414
    Examples of SPT Threshold Policy Configuration 414
Configuring Any-Source Multicast for Draft-Rosen VPNs 415
    Configuring the VPN 416
    Configuring PIM Connectivity Between the Provider and PE Routers 417
    Configuring Multicast Connectivity on the CE Routers 417
    Configuring Multicast Connectivity for the VPN on the PE Router 418
    Configuring the Routing Group 418

Chapter 42 **PIM Configuration Examples** **421**

Example: Configuring PIM Dense Mode 421
Example: Configuring PIM Sparse-Dense Mode 422
Example: Configuring PIM Sparse Mode 422
Example: Configuring PIM Anycast with MSDP 423
Example: Configuring PIM BSR Filters 427
Example: Filtering Incoming PIM Join Messages 427

Example: Configuring Register Message Filters on PIM RPs and DRs 428
Example: Configuring Externally Facing PIM Border Routers 430
Example: Tracing PIM Protocol Traffic 430
Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain 431
Example: Configuring PIM Dense Mode over Layer 3 VPNs 431
Example: Configuring PIM Sparse-Dense Mode over Layer 3 VPNs 433
Example: Configuring PIM Sparse Mode over Layer 3 VPNs 436

Chapter 43 Summary of PIM Configuration Statements 445

address (Anycast RPs) 445
address (Local RPs) 446
address (Static RPs) 447
algorithm 448
anycast-pim 449
assert-timeout 450
authentication 451
auto-rp 452
bfd-liveness-detection 453
bootstrap 454
bootstrap-export 454
bootstrap-import 455
bootstrap-priority 456
dense-groups 456
disable (PIM Graceful Restart) 457
disable (PIM) 458
dr-election-on-p2p 459
dr-register-policy 459
embedded-rp 460
export (Bootstrap) 460
export (PIM) 461
family (Bootstrap) 461
family (Disable PIM) 462
family (Local RP) 463
graceful-restart 464
group-ranges 465
hello-interval 466
hold-time 467
import (Bootstrap) 467
import (PIM) 468
infinity 468
interface 469
join-prune-timeout 470
join-load-balance 470
key-chain 471
local 472
local-address 473
loose-check 473
mapping-agent-election 474

maximum-rps 474
minimum-interval 475
minimum-receive-interval 475
mode 476
multiplier 477
neighbor-policy 477
pim 478
priority (Bootstrap) 480
priority (PIM Interfaces) 481
priority (PIM RPs) 481
restart-duration 482
rib-group 482
rp 483
rp-register-policy 484
rp-set 485
spt-threshold 486
static 487
traceoptions 488
version (BFD) 491
version (PIM) 492
vpn-group-address 492

Part 12 **MSDP**

Chapter 44 **MSDP Introduction** **495**

MSDP Overview 495

Chapter 45 **MSDP Configuration Guidelines** **497**

Configuring MSDP 497
Minimum MSDP Configuration 499
Enabling MSDP 499
Configuring MSDP Peers 500
Configuring MSDP Groups 501
Configuring MSDP Mesh Groups 502
Configuring the MSDP Authentication Key 503
Configuring MSDP Routing Policy 504
Configuring Multiple Rendezvous Points in a Domain 506
    Example: Configuring a Router to Use Anycast RP 506
Configuring MSDP Data Encapsulation 507
Configuring the MSDP Active Source Limit 508
    Configuring Global, Group, and Peer Active Source Limit 509
    Configuring Per-Source Active Source Limit 510
Configuring a Default MSDP Peer 510

Disabling MSDP 511
Tracing MSDP Protocol Traffic 512
Example: Tracing MSDP Protocol Traffic 513
Example: Configuring MSDP 513

Chapter 46 **Summary of MSDP Configuration Statements** **515**

active-source-limit 516
authentication-key 517
data-encapsulation 518
default-peer 519
disable 520
export 521
group 522
import 523
local-address 524
maximum 525
mode 526
msdp 527
peer 529
rib-group 530
source 531
threshold 532
traceoptions 533

Part 13 **Index**

Index 539
Index of Statements and Commands 551

List of Figures

Part 1 **Introduction to Multicast**

Chapter 1 **Multicast Overview** **3**

Figure 1: Multicast Terminology in an IP Network 6
Figure 2: Multicast Routers and the RPF Check 11

Chapter 2 **IP Multicast Overview** **17**

Figure 3: Converting MAC Addresses to Multicast Addresses 22

Part 4 **MLD**

Chapter 8 **MLD Overview Introduction** **95**

Figure 4: Routers Start Up on a Subnet 96
Figure 5: Querier Router Is Determined 96
Figure 6: General Query Message Is Issued 96
Figure 7: Reports Are Received by the Querier Router 97
Figure 8: Host Has No Interested Receivers and Sends a Done Message to Router 97
Figure 9: Host Address Timer Expires and Address Is Removed from Multicast Address List 97

Part 6 **PGM**

Chapter 14 **PGM Overview** **149**

Figure 10: PGM Architecture and General Operation 153

Part 8 **Multicast Routing Options**

Chapter 22 **Source-Specific Multicast Configuration Guidelines** **203**

Figure 11: Receiver Sends Messages to Join Group G and Source S 204
Figure 12: Router 3 (Last-Hop Router) Joins the Source Tree 205
Figure 13: The (S,G) State Is Built Between the Source and the Receiver 205
Figure 14: Network on Which to Configure PIM SSM 207

Part 11 **PIM**

Chapter 38 **PIM Overview** **341**

Figure 15: Rendezvous Point as Part of the RPT and SPT 344
Figure 16: Building an RPT Between RP and Receiver 347
Figure 17: PIM Register Message and PIM Join Message Exchanged 348
Figure 18: Traffic Sent from the Source to the RP Router 349
Figure 19: Traffic Sent from the RP Router Toward the Receiver 350

Figure 20: Receiver DR Sends a PIM Join Message to the Source 351
Figure 21: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router 351
Figure 22: RP Router Receives PIM Prune Message 352
Figure 23: RP Router Sends a PIM Prune Message to the Source DR 353
Figure 24: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router 353
Figure 25: Receiver Announces Desire to Join Group G and Source S 355
Figure 26: Router 3 (Last-Hop Router) Joins the Source Tree 355
Figure 27: The (S,G) State Is Built Between the Source and the Receiver 355
Figure 28: Multicast Traffic Flooded from the Source Using PIM Dense Mode 356
Figure 29: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic 357
Figure 30: Source and Receiver Sites in an MVPN 361
Figure 31: Adding a Receiver to an MVPN Source Site Using MBGP 362
Figure 32: Extracting the Embedded RP IPv6 Address 367

Chapter 41 PIM Sparse Mode Configuration Guidelines 393

Figure 33: Configuring the VPN 416
Figure 34: Multicast Configuration on the Provider Network 417
Figure 35: Multicast Connectivity on the CE Routers 418
Figure 36: Multicast Connectivity for the VPN 418

Chapter 42 PIM Configuration Examples 421

Figure 37: Customer Edge and Service Provider Networks 436

Part 12 MSDP

Chapter 45 MSDP Configuration Guidelines 497

Figure 38: Source-Active Message Flooding 503

List of Tables

About This Guide **xxxi**

Table 1: Notice Icons xxxv

Table 2: Text and Syntax Conventions xxxv

Part 1 **Introduction to Multicast**

Chapter 1 **Multicast Overview** **3**

Table 3: Multicast Routing Protocols Compared 16

Part 2 **IP Multicast Configuration**

Chapter 3 **Introduction to PIM** **27**

Table 4: ASM and SSM Terminology 31

Part 3 **IGMP**

Chapter 6 **IGMP Configuration Guidelines** **55**

Table 5: IGMP Event Messages 70

Part 4 **MLD**

Chapter 9 **MLD Configuration Guidelines** **99**

Table 6: MLD Event Messages 112

Part 11 **PIM**

Chapter 38 **PIM Overview** **341**

Table 7: Tunnel PIC Requirements for IPv4 and IPv6 Multicast 363

Chapter 39 **PIM Mode-Independent Configuration Guidelines** **369**

Table 8: PIM Join Filter Match Conditions 377

Chapter 41 **PIM Sparse Mode Configuration Guidelines** **393**

Table 9: Local RP and Auto-RP Message Types 403

Part 12 **MSDP**

Chapter 45 **MSDP Configuration Guidelines** **497**

Table 10: Source-Active Message Flooding Explanation 503

Table 11: MSDP Source-Active Message Filter Match Conditions 505

About This Guide

This preface provides the following guidelines for using the *JUNOS® Software Multicast Protocols Configuration Guide*:

- JUNOS Documentation and Release Notes on page xxxi
- Objectives on page xxxii
- Audience on page xxxii
- Supported Platforms on page xxxii
- Using the Indexes on page xxxiii
- Using the Examples in This Manual on page xxxiii
- Documentation Conventions on page xxxiv
- Documentation Feedback on page xxxvi
- Requesting Technical Support on page xxxvi

JUNOS Documentation and Release Notes

For a list of related JUNOS documentation, see http://www.juniper.net/techpubs/software/junos/.

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

Juniper Networks supports a technical book program to publish books by Juniper Networks engineers and subject matter experts with book publishers around the world. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration using JUNOS Software and Juniper Networks devices. In addition, the Juniper Networks Technical Library, published in conjunction with O'Reilly Media, explores improving network security, reliability, and availability using JUNOS configuration techniques. All the books are for sale at technical bookstores and book outlets around the world. The current list can be viewed at http://www.juniper.net/books .

Objectives

This guide provides an overview of the multicast protocols for the JUNOS Software and describes how to configure multicast protocols on the router.

NOTE: For additional information about JUNOS Software—either corrections to or information that might have been omitted from this guide—see the software release notes at http://www.juniper.net/.

Audience

This guide is designed for network administrators who are configuring and monitoring a Juniper Networks M Series, MX Series, T Series, EX Series, or J Series router or switch.

To use this guide, you need a broad understanding of networks in general, the Internet in particular, networking principles, and network configuration. You must also be familiar with one or more of the following Internet routing protocols:

- Border Gateway Protocol (BGP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Internet Control Message Protocol (ICMP) router discovery
- Internet Group Management Protocol (IGMP)
- Multiprotocol Label Switching (MPLS)
- Open Shortest Path First (OSPF)
- Protocol-Independent Multicast (PIM)
- Resource Reservation Protocol (RSVP)
- Routing Information Protocol (RIP)
- Simple Network Management Protocol (SNMP)

Personnel operating the equipment must be trained and competent; must not conduct themselves in a careless, willfully negligent, or hostile manner; and must abide by the instructions provided by the documentation.

Supported Platforms

For the features described in this manual, JUNOS Software currently supports the following platforms:

- J Series
- M Series

- MX Series
- T Series
- EX Series

Using the Indexes

This reference contains two indexes: a complete index that includes topic entries, and an index of statements and commands only.

In the index of statements and commands, an entry refers to a statement summary section only. In the complete index, the entry for a configuration statement or command contains at least two parts:

- The primary entry refers to the statement summary section.
- The secondary entry, *usage guidelines,* refers to the section in a configuration guidelines chapter that describes how to use the statement or command.

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

```
system {
   scripts {
      commit {
         file ex-script.xsl;
      }
   }
}
interfaces {
   fxp0 {
```

```
            disable;
            unit 0 {
                family inet {
                    address 10.0.0.1/24;
                }
            }
        }
    }
```

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

Table 1 on page xxxv 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.

Table 2 on page xxxv 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**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host> show chassis alarms` `No alarms currently active`
Italic text like this	■ Introduces important new terms. ■ Identifies book names. ■ Identifies RFC and Internet draft titles.	■ A policy *term* is a named structure that defines match conditions and actions. ■ *JUNOS System Basics Configuration Guide* ■ RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	■ To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. ■ The console port is labeled CONSOLE.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric *metric*>;

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

Convention	Description	Examples
| (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)	Enclose a variable for which you can substitute one or more values.	community name members [*community-ids*]

Indention and braces ({ })	Identify 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.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	■ In the Logical Interfaces box, select **All Interfaces**. ■ To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web 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 send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support

contract, or are covered under warranty, and need postsales 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/customers/support/downloads/710059.pdf .
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/ .
- JTAC Hours of Operation —The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

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

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

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

Part 1
Introduction to Multicast

- Multicast Overview on page 3
- IP Multicast Overview on page 17

Chapter 1

Multicast Overview

This chapter discusses the following topics:

- Comparing Multicast to Unicast on page 3
- IP Multicast Uses on page 5
- IP Multicast Terminology on page 6
- Multicast Leaf and Branch Terminology on page 7
- Dense and Sparse Modes for Multicast Networks on page 7
- IP Multicast Building Blocks Overview on page 8
- IP Multicast Addressing on page 8
- Multicast Administrative Scoping on page 9
- Multicast Interface Lists on page 9
- Multicast Reverse-Path Forwarding on page 10
- Multicast Shortest-Path Tree on page 11
- Multicast Rendezvous Points, Shared Trees, and the Rendezvous-Point Trees on page 12
- Multicast RPF Checks on page 13
- Populating the Multicast RPF Table on page 14
- Protocols for Multicast on page 14
- Multicast Group Membership Protocols on page 14
- Multicast Routing Protocols on page 15

Comparing Multicast to Unicast

The JUNOS Software routing protocol process supports a wide variety of routing protocols. These routing protocols carry network information among routers 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 Web site server, to a single destination such as a client PC. This 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 public 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 routers 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 routers 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 routers. Multicast routers distribute the multicast traffic across the network from source to destinations. The multicast router 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 Web site replication, and distributed interactive simulation (DIS) such as war games 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 routers 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 routers replicate the packets and deliver the packets to the proper receivers, but only the replication role is a new one for routers. 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 routers and networks. Figure 1 on page 6 shows a general view of some of the terms commonly used in an IP multicast network.

In a multicast network, the key component is the *router*, able to replicate packets and therefore multicast-capable. The routers 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 router 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 router to receive multicast packets. The interface on the router 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 router, where *N* is the number of logical interfaces on the router. To prevent looping, the upstream interface must never receive copies of downstream multicast packets.

Figure 1: Multicast Terminology in an IP Network

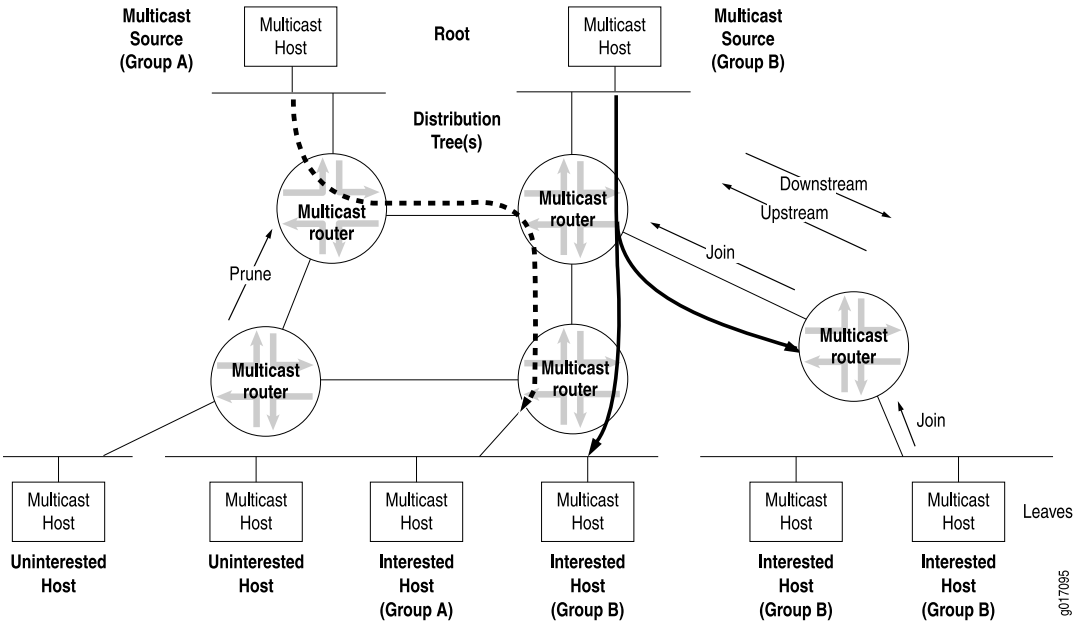

Routing loops are disastrous in multicast networks because of the risk of repeatedly replicated packets. One of the complexities of modern multicast routing protocols is the need to avoid routing loops, packet by packet, much more rigorously than in unicast routing protocols.

Multicast terminology includes two more complex concepts:

- Multicast Leaf and Branch Terminology on page 7
- Dense and Sparse Modes for Multicast Networks on page 7

Multicast Leaf and Branch Terminology

Each subnetwork with hosts on the router that has at least one interested receiver is a *leaf* on the distribution tree. Routers 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 now sent out on the interface. The number of leaves on a particular interface does not affect the router. The action is the same for one leaf or a hundred.

When a branch contains no leaves because there are no interested hosts on the router 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 router, 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.

Dense and Sparse Modes for Multicast Networks

The actions of receivers suggest two basic strategies for protocols to handle joining and pruning branches among a collection of multicast routers:

- Dense-mode multicast—The assumption could be made 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.
- Sparse-mode multicast—Alternatively, the assumption could be made that very few of the possible receivers want packets from this 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. This is the *sparse mode* of multicast operation. 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.

Some multicast routing protocols, especially older ones, support only dense-mode operation, which makes them inappropriate for use on the public Internet. Others allow sparse mode as well. If *sparse-dense mode* is supported, the multicast routing protocol allows some multicast groups to be sparse and other groups to be dense.

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

IP Multicast Building Blocks Overview

Implementing an IP multicast network is made simpler by using a number of standardized building blocks. These IP multicast building blocks include a special IP multicast address space, administrative scoping to prevent large-scale routing loops, upstream and downstream interface lists, reverse path forwarding (RPF) to prevent small-scale routing loops, a shortest-path tree (SPT) algorithm to build a minimal distribution tree, and a rendezvous point (RP) and associated rendezvous-point tree (RPT) to allow sparse mode receivers to find sources.

There are six major IP multicast building blocks:

- IP Multicast Addressing on page 8
- Multicast Administrative Scoping on page 9
- Multicast Interface Lists on page 9
- Multicast Reverse-Path Forwarding on page 10
- Multicast Shortest-Path Tree on page 11
- Multicast Rendezvous Points, Shared Trees, and the Rendezvous-Point Trees on page 12

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 are concerned more with content than with physical devices. Receivers are not assigned their own multicast addresses, but need

to know only 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 Administrative Scoping

Routing loops must be avoided in IP multicast networks. Because multicast routers must replicate packets for each downstream branch, not only do looping packets not arrive at a destination, but each pass around the loop multiplies the number of looping packets, eventually overwhelming the network.

Scoping limits the routers and interfaces that can be used to forward a multicast packet. Scoping can use the time-to-live (TTL) field in the IP packet header, but TTL scoping depends on intimate knowledge of the network topology by the network administrator. This topology can change as links fail and are restored, making TTL scoping a poor solution for multicast.

Multicast scoping is *administrative* in the sense that a range of multicast addresses is reserved for scoping purposes, as described in RFC 2365. Routers at the boundary must be able to filter multicast packets and make sure the packets do not stray beyond the established limit.

Administrative scoping is much better than TTL scoping, but in many cases the dropping of administratively scoped packets is still determined by the network administrator. For example, the multicast address range 239/8 is defined in RFC 2365 as administratively scoped, and packets using this range are not to be forwarded beyond a network “boundary,” usually a routing domain. But only the network administrator knows where the border routers are and can implement the scoping correctly.

Multicast groups used by unicast routing protocols, such as 224.0.0.5 for all Open Shortest Path First (OSPF) routers, are administratively scoped for that LAN only. This scoping allows the same multicast address to be used without conflict on every LAN running OSPF.

Multicast Interface Lists

To avoid multicast routing loops, every multicast router 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.

Routers closely monitor the status of the incoming and outgoing interfaces, a process that determines the *multicast forwarding state*. A router with a multicast forwarding state for a particular multicast group is essentially “turned on” for that group's content. Interfaces on the router'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 router 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 router could use (*,224.1.1.2) to represent the state of a router forwarding traffic from both sources to the group.

For more information about the use of multicast forwarding state notations in different types of distribution trees, see “Multicast Rendezvous Points, Shared Trees, and the Rendezvous-Point Trees” on page 12. For more information about the use of multicast notations in different multicast routing protocols, see “Multicast Group Membership Protocols” on page 14 and “Multicast Routing Protocols” on page 15.

Multicast Reverse-Path Forwarding

Unicast forwarding decisions are typically based on the destination address of the packet arriving at a router. The unicast routing table is organized by destination subnet and mainly set up to forward the packet toward the destination.

In multicast, the router forwards the packet away from the source to make progress along the distribution tree and prevent routing loops. The router's multicast forwarding state runs more logically by organizing tables based on the reverse path, from the receiver back to the root of the distribution tree. This process is known as *reverse-path forwarding* (*RPF*).

The router adds a branch to a distribution tree depending on whether the request for traffic from a multicast group passes the reverse-path-forwarding check (RPF check). Every multicast packet received must pass an RPF check before it is eligible to be replicated or forwarded on any interface.

The RPF check is essential for every router's multicast implementation. When a multicast packet is received on an interface, the router interprets the source address in the multicast IP packet as the destination address for a unicast IP packet. The source multicast address is found in the unicast routing table, and the outgoing interface is determined. If the outgoing interface found in the unicast routing table is the same as the 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.

Figure 2 on page 11 shows how multicast routers can use the unicast routing table to perform an RPF check and how the results obtained at each router determine where join messages are sent.

Figure 2: Multicast Routers and the RPF Check

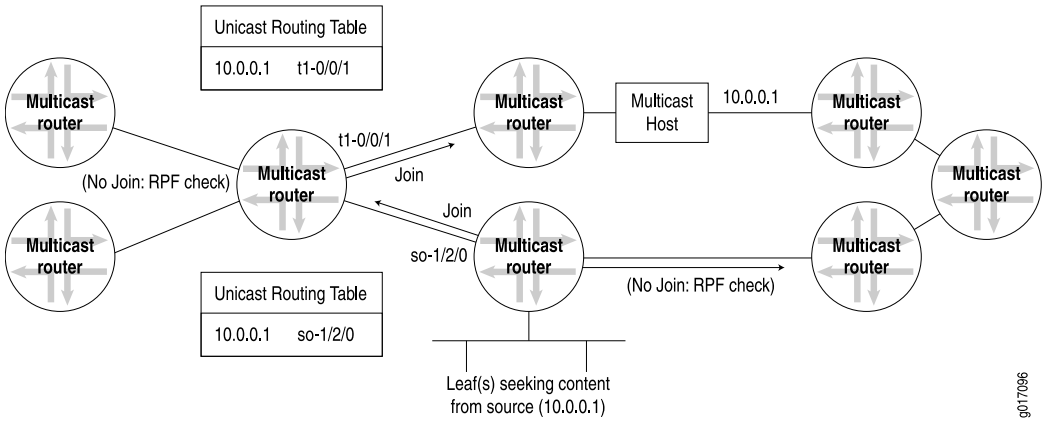

Routers can build and maintain separate tables for RPF purposes. The router must have some way to determine its *RPF interface* for the group, which is the interface topologically closest to the root. For greatest efficiency, the distribution tree follows the shortest-path tree topology. The RPF check helps to construct this tree.

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 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 wishes 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 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 the RPF interface, building the SPT as it goes. The process stops when the join message:

- Reaches the router directly connected to the host that is the source, or
- 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 could 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 "Multicast Rendezvous Points, Shared Trees, and the Rendezvous-Point Trees" on page 12.

Multicast Rendezvous Points, Shared Trees, and the 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 knows 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 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 knows 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 wishes to join the shared tree for that group. This message is a (*,G) join message because S is not known, only the RP, 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 OIL for the group and also performs an RPF check on the RP address. The upstream router then sends a (*,G) join message out 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 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 that group
- A router along the RPT that already has a multicast forwarding state for this group

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 "Multicast RPF Checks" on page 13.

Multicast RPF Checks

The RPF table plays the key role in the multicast router. The RPF table is consulted for every RPF check, which is performed at intervals on multicast packets entering the multicast router. Distribution trees of all types rely on the RPF table to form properly, and the multicast forwarding state also depends on the RPF table.

RPF checks are performed only on unicast addresses to find the upstream interface for the multicast source or RP.

The routing table used for RPF checks can be the same routing table used to forward unicast IP packets, or it can be a separate routing table used only for multicast RPF checks. In either case, the RPF table contains only unicast routes, because the RPF check is performed on the source address of the multicast packet, not the multicast group destination address, and a multicast address is forbidden from appearing in the source address field of an IP packet header. The unicast address can be used for RPF checks because there is only one source host for a particular stream of IP multicast content for a multicast group address, although the same content could be available from multiple sources.

Populating the Multicast RPF Table

If the same routing table used to forward unicast packets is also used for the RPF checks, the routing table is populated and maintained by the traditional unicast routing protocols such as Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), OSPF, and Routing Information Protocol (RIP). If a dedicated multicast RPF table is used, this table must be populated by some other method. Some multicast routing protocols (such as the Distance Vector Multicast Routing Protocol [DVMRP]) essentially duplicate the operation of a unicast routing protocol and populate a dedicated RPF table. Others, such as PIM, do not duplicate routing protocol functions and must rely on some other routing protocol to set up this table, which is why PIM is *protocol independent*.

Some traditional routing protocols such as BGP and IS-IS now have extensions to differentiate between different sets of routing information sent between routers for unicast and multicast. For example, there is multiprotocol BGP (MBGP) and multitopology routing in IS-IS (M-ISIS). IS-IS routes can be added to the RPF table even when special features such as traffic engineering and “shortcuts” are turned on. Multicast Open Shortest Path First (MOSPF) also extends OSPF for multicast use, but goes further than MBGP or M-ISIS and makes MOSPF into a complete multicast routing protocol on its own. When these routing protocols are used, routes can be tagged as multicast RPF routers and used by the receiving router differently than the unicast routing information.

Using the main unicast routing table for RPF checks provides simplicity. A dedicated routing table for RPF checks allows a network administrator to set up separate paths and routing policies for unicast and multicast traffic, allowing the multicast network to function more independently of the unicast network.

Protocols for Multicast

The protocols used among a collection of multicast-capable IP routers fall into two major categories:

- Multicast group membership protocols that are used between host and router (and host to host)
- Multicast routing protocols that are used between routers

Multicast Group Membership Protocols

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

There is one standard IP multicast group membership protocol: the Internet Group Management Protocol (IGMP). However, IGMP has several versions that are supported by hosts and routers. There are currently three versions of IGMP:

- IGMPv1—The original protocol defined in RFC 1112. An explicit join message is sent to the router, but a timeout is used to determine when hosts leave a group. This process wastes processing cycles on the router, especially on older or smaller routers.
- IGMPv2—Defined in RFC 2236. Among other features, IGMPv2 adds an explicit leave message to the join message so that routers can more easily determine when a group has no interested listeners on a LAN.
- IGMPv3—Defined in RFC 3376. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or *source-specific multicast* (*SSM*). (RFC 1112 supported both many-to-many and one-to-many multicast, but one-to-many is considered the more viable model for the Internet at large.)

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

Multicast Routing Protocols

Multicast routing protocols enable a collection of multicast routers 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.

There are five multicast routing protocols:

- 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).
- MOSPF—Extends OSPF for multicast use, but only for dense mode. However, MOSPF has an explicit join message, so routers 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).
- PIM dense mode—This is PIM operating in dense mode (PIM DM), but the differences from PIM sparse mode are profound enough to consider the two modes separately. PIM also supports sparse-dense mode, with mixed sparse and dense groups, but there is no special notation for that operational mode. In contrast to DVMRP and MOSPF, PIM dense mode allows a router to use any unicast routing protocol and performs RPF checks using the unicast routing table. PIM dense mode has an implicit join message, so routers 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 sparse mode—Allows a router to use any unicast routing protocol and performs RPF checks using the unicast routing table. However, PIM sparse mode has an *explicit* join message, so routers determine where the interested receivers are and send join messages upstream to their neighbors, building trees from receivers to RP. PIM sparse mode uses an RP router as the initial source of multicast group traffic and therefore builds distribution trees in the form (*,G), as do all sparse-mode protocols. However, 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.
- 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.

The differences among the five multicast routing protocols are summarized in Table 3 on page 16.

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
CBT	No	Yes	No	Yes	No	Yes

It is important to realize that retransmissions due to a high bit-error rate on a link or overloaded router can make multicast as inefficient as repeated unicast. Therefore, there is a trade-off in many multicast applications regarding the session support provided by 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.

Chapter 2

IP Multicast Overview

This chapter discusses the following topics:

- JUNOS Software Multicast Standards on page 17
- IP Multicast Specifications on page 18
- Multicast Overview on page 19
- Multicast Addresses on page 20
- Multicast Redundancy on page 20
- T Series Router Multicast Performance on page 21
- Layer 2 Frames and IPv4 Multicast Addresses on page 21
- Multicast Snooping Overview on page 23

JUNOS Software Multicast Standards

The JUNOS Software implements the following protocols to support IP multicast routing:

- Internet Group Management Protocol (IGMP), versions 1 and 2—Used to learn whether group members are present, for IP version 4 (IPv4) routers.
- Multicast Listener Discovery (MLD), versions 1 and 2—Used to learn whether group members are present, for IP version 6 (IPv6) routers.
- Distance Vector Multicast Routing Protocol (DVMRP)—Dense-mode multicast routing protocol.
- Protocol Independent Multicast (PIM)—Multicast routing protocol that routes to multicast groups that might span wide-area and interdomain internetworks. Both dense mode and sparse mode are supported.
- Multicast Source Discovery Protocol (MSDP)—Multicast routing protocol that discovers active sources of multicast messages. PIM sparse mode uses these sources.
- Session Announcement Protocol (SAP) and Session Description Protocol (SDP)—Handle conference session announcements.

IP Multicast Specifications

The protocols related to IP multicast are defined in the following documents:

- RFC 1112, *Host Extensions for IP Multicasting* (defines IGMP Version 1)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 2327, *SDP: Session Description Protocol*
- RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*
- RFC 2365, *Administratively Scoped IP Multicast*
- RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*
- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 2974, *Session Announcement Protocol*
- RFC 3208, *PGM Reliable Transport Protocol Specification*
- RFC 3376, *Internet Group Management Protocol, Version 3* (source-specific multicast [SSM] include and exclude mode)
- RFC 3446, *Anycast Rendezvous Point (RP) Mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*
- RFC 3569, *An Overview of Source-Specific Multicast (SSM)*
- RFC 3590, *Source Address Selection for Multicast Listener Discovery Protocol* (SSM include and exclude mode)
- RFC 3618, *Multicast Source Discovery Protocol (MSDP)*
- RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2)*
- RFC 3973, *Protocol Independent Multicast—Dense Mode (PIM-DM)*
- RFC 4364, *BGP/MPLS IP VPNs*
- RFC 4601, *Protocol Independent Multicast—Sparse Mode (PIM-SM)*
- Internet draft draft-ietf-pim-sm-bsr-05.txt, *Bootstrap Router (BSR) Mechanism for PIM Sparse Mode* (expired August 2005) (no support for draft's scoping mechanism)
- Internet draft draft-ietf-idmr-dvmrp-v3-11.txt, *Distance Vector Multicast Routing Protocol* (expired April 2004)
- Internet draft draft-raggarwa-l3vpn-2547-mvpn-00.txt, *Base Specification for Multicast in BGP/MPLS VPNs,* Section 2 (expired December 2004)
- Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP VPNs*, data MDTs only (expired April 2004)
- Internet draft draft-ietf-ssm-arch-06.txt, *Source-Specific Multicast for IP* (expired March 2005)
- Internet draft draft-ietf-mboned-ssm232-08.txt, *Source-Specific Protocol Independent Multicast in 232/8* (expired September 2004)

- Internet draft draft-holbrook-idmr-igmpv3-ssm-07.txt, *Using IGMPv3 and MLDv2 for Source-Specific Multicast* (expired December 2004)
- Internet draft draft-ietf-l3vpn-2547bis-mcast-02.txt, *Multicast in MPLS/BGP IP VPNs* (expired September 2007)
- Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-03.txt, *BGP Encodings for Multicast in MPLS/BGP IP VPNs* (expired April 2007)
- Internet draft draft-raggarwa-l3vpn-2547-mvpn-00.txt, *Base Specification for Multicast in BGP/MPLS VPNs* (expired December 2004)

NOTE: The implementation of the Distance Vector Multicast Routing Protocol (DVMRP) is based on a series of expired Internet drafts, as are all current implements of DVMRP. None are based on RFC 1075, *Distance Vector Multicast Routing Protocol*.

To access Internet RFCs and drafts, go to the IETF Web site at http://www.ietf.org.

Multicast Overview

IPv4 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.

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 and 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.

The Internet Multicast Backbone (MBone) is an interconnected set of subnetworks and routers that support the delivery of IP multicast traffic. The MBone is a virtual network that is layered on top of sections of the physical Internet. The MBone is composed of islands of multicast routing capability that are connected to other islands by virtual point-to-point links called tunnels. The tunnels allow multicast traffic to pass undisturbed through the parts of the Internet that are not multicast-capable.

Because the MBone and the Internet have different topologies, multicast routers execute a separate routing protocol to decide how to forward multicast packets.

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.

Multicast Redundancy

The JUNOS Software supports nondisruptive, graceful Routing Engine switchover and graceful restart for some routing protocols in the case of Routing Engine or routing process failure. You can configure many routing protocols to continue to forward packets during the Routing Engine switchover or routing process restart. To support graceful Routing Engine switchover, the router must have two Routing Engines installed and be configured properly. For more information about graceful Routing Engine switchover, see the *JUNOS High Availability Configuration Guide*.

Graceful restart of routing protocol processes is required for graceful switchover. By default, graceful restart is disabled and must be enabled at the [edit routing-options graceful-restart] hierarchy level.

PIM sparse mode uses a mechanism called a *generation identifier* to indicate the need for graceful restart. Generation identifiers are included by default in PIM hello messages, as specified in the Internet draft draft-ietf-pim-sm-v2-new-10.txt. An initial generation identifier is created by each PIM neighbor to establish device capabilities. When one of the PIM neighbors restarts, it sends a new generation identifier to its neighbors. All neighbors that support graceful restart and are connected by point-to-point links assist by sending multicast updates to the restarting neighbor.

The restart phase is completed either when the PIM state becomes stable or when the restart interval timer expires. If the neighbors do not support graceful restart or if they connect to each other using multipoint interfaces, the restarting router uses a restart interval timer to define the restart period.

Graceful restart is compatible with the use of all multicast protocols. However, only PIM benefits from this feature. The router does not forward multicast packets for protocols other than PIM during graceful restart or switchover, because all other multicast protocols must completely restart after a routing process failure.

To configure graceful restart for PIM, include the graceful-restart statement at the [edit routing-options] hierarchy level, and include the pim statement at the [edit protocols] hierarchy level:

```
[edit routing-options]
graceful-restart;
[edit protocols]
pim {...};
```

For more information about graceful restart for PIM sparse mode, see "Configuring PIM Sparse Mode Graceful Restart" on page 395 and the *JUNOS Feature Guide*. For more information about graceful restart and other routing protocols, see the *JUNOS Routing Protocols Configuration Guide*.

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.

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 routers, 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 a 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 three 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 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 procedure is shown in Figure 3 on page 22.

Figure 3: Converting MAC Addresses to Multicast Addresses

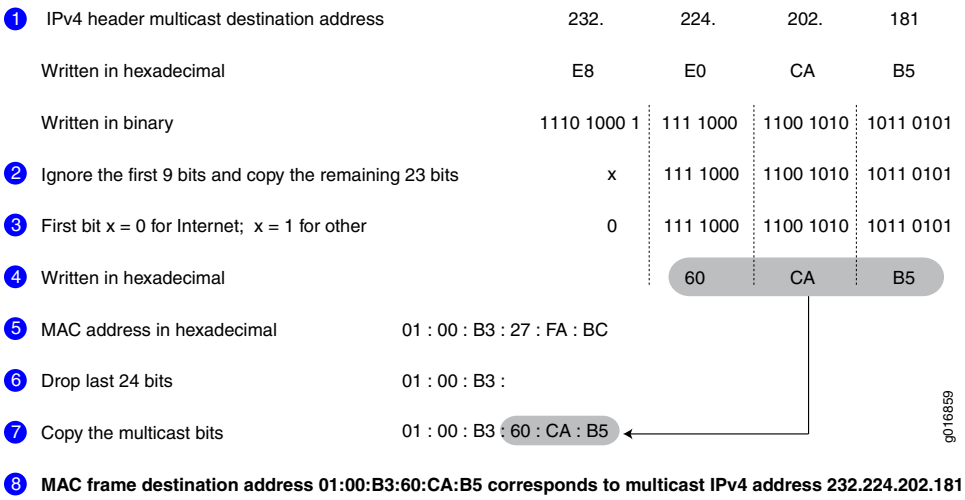

Note that this process means that there are 32 (2^5) IPv4 multicast addresses that could map to the same MAC multicast addresses. For example, multicast IPv4 addresses 224.8.7.6 and 229.136.7.6 translate to the same MAC address (0x01-00-5E-08-07-06). This is a real concern, and because the host could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

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 Snooping Overview

Network devices such as routers operate mainly at the packet level, or Layer 3. Other network devices such as bridges or LAN switches operate mainly at the frame level, or Layer 2. Multicasting functions mainly at the packet level, Layer 3, but there is a way to map Layer 3 IP multicast group addresses to Layer 2 MAC multicast group addresses at the frame level. For more information about this translation procedure, see “Layer 2 Frames and IPv4 Multicast Addresses” on page 21.

Routers are able to handle both Layer 2 and Layer 3 addressing information because the frame and its addresses must be processed to access the encapsulated packet inside. It is easy for routers to run Layer 3 multicast protocols such as PIM or IGMP and determine where to forward multicast content or when a host on an interface joins or leaves a group. However, bridges and LAN switches, as Layer 2 devices, are not supposed to have access to the multicast information inside the packets their frames carry.

How then are bridges and other Layer 2 devices to determine when a device on an interface joins or leaves a multicast tree, or whether a host on an attached LAN is interested in receiving the content of a particular multicast group or not?

The answer is for the Layer 2 device to implement a series of procedures known as *multicast snooping*. Multicast snooping is a general term and applies to the process of a Layer 2 device “snooping” at the Layer 3 packet content to determine which actions are taken to process or forward a frame. There are more specific forms of snooping, such as IGMP snooping or PIM snooping. In all cases, snooping involves a device configured to function at Layer 2 having access to normally “forbidden” Layer 3 (packet) information. However, snooping makes multicasting more efficient in these devices.

Part 2
IP Multicast Configuration

- Introduction to PIM on page 27
- Complete Multicast Configuration Statements on page 33

Chapter 3

Introduction to PIM

This chapter provides an overview of the features and capabilities of PIM:

- PIM Background on page 27
- Basic PIM Network Components on page 28
- PIM Modes of Operation on page 28
- PIM Dense Mode Overview on page 29
- PIM Sparse Mode Overview on page 29
- PIM Source-Specific Mode Overview on page 30
- Mixing PIM Sparse and Dense Modes on page 31

PIM Background

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 problems 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 or 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.

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 everything 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. So 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).

PIM Modes of Operation

PIM operates in two basic modes: sparse mode and dense mode. In addition, PIM can operate in sparse-dense mode, with some multicast groups configured as dense

mode (flood-and-prune, [S,G] state) and others 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.

PIM Dense Mode Overview

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 need not introduce and manage a separate routing protocol just for RPF checks. Unicast routing protocols extended as multiprotocol BGP (MBGP) and Multitopology Routing in Intermediate System-to-Intermediate System (M-ISIS) were later employed to build special tables to perform RPF checks, but PIM dense mode does not require them.

PIM dense mode can use the unicast routing table populated by Open Shortest Path First (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-ISIS when performing RPF checks.

PIM Sparse Mode Overview

These are the major characteristics of PIM sparse mode:

- 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 then get content directly from the source.
- This transitional aspect of PIM sparse mode from shared to source-based tree is one of the major attractions of PIM. This feature 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 each other.
- 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 CBT.

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

PIM Source-Specific Mode Overview

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 knows 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 are responsible for knowing 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 makes perfect sense. However, by ignoring the many-to-many model and focusing attention on the one-to-many 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 4 on page 31.

Table 4: 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.

Mixing PIM Sparse and Dense Modes

It is possible to mix PIM dense mode, PIM sparse mode, and PIM SSM on the same network, the same router, 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.

A multicast router employing sparse-dense mode is a good example of mixing PIM modes on the same network or router or interface. Dense modes are easy to support because of the flooding, but the scaling issues make dense modes inappropriate for Internet use beyond very restricted uses.

PIM sparse mode was capable of forming 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.

Chapter 4

Complete Multicast Configuration Statements

This chapter is organized as follows:

- Multicast Configuration Statements Overview on page 33
- [edit protocols] Hierarchy Level on page 33
- [edit routing-instances] Hierarchy Level on page 38
- [edit routing-options] Hierarchy Level on page 40
- [edit logical-systems protocols] Hierarchy Level on page 41
- [edit logical-systems routing-instances] Hierarchy Level on page 45
- [edit logical-systems routing-options] Hierarchy Level on page 48
- [edit bridge-domains bridge-domain-name protocols] Hierarchy Level on page 49

Multicast Configuration Statements Overview

This chapter shows the complete configuration statement hierarchy for the portions of the configuration discussed in this manual, listing all possible configuration statements and showing their level in the configuration hierarchy. When you are configuring the JUNOS Software, your current hierarchy level is shown in the banner on the line preceding the user@host# prompt.

For the complete configuration statement hierarchy, see the *JUNOS Hierarchy and RFC Reference*.

[edit protocols] Hierarchy Level

```
[edit]
protocols {
```

Distance Vector Multicast Routing Protocol (DVMRP)

```
dvmrp {
   disable;
   export [ policy-names ];
   import [ policy-names ];
   interface interface-name {
      disable;
      hold-time seconds;
      metric metric;
      mode (forwarding | unicast-routing);
```

```
    }
    rib-group group-name;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Internet Group Management Protocol (IGMP)

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

Multicast Listener Discovery (MLD)

```
mld {
    accounting;
    interface interface-name {
        disable;
        (accounting | no-accounting);
        group-policy;
        immediate-leave;
        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>;
  }
}
```

Multicast Source Discovery Protocol (MSDP)

```
msdp {
  disable;
  active-source-limit {
    maximum number;
    threshold number;
  }
  data-encapsulation (disable | enable);
  export [ policy-names ];
  group group-name {
    ... group-configuration ...
  }
  import [ policy-names ];
  local-address address;
  peer address {
    ... peer-configuration ...
  }
  rib-group group-name;
  source ip-prefix</prefix-length> {
    active-source-limit {
      maximum number;
      threshold number;
    }
  }
  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>;
        }
    }
}
```

Pragmatic General Multicast (PGM)

```
pgm {
    traceoptions {
        flag flag <flag-modifier> <disable>;
        }
    }
}
```

Protocol Independent Multicast (PIM)

```
pim {
    disable;
    assert-timeout seconds;
    dense-groups {
        addresses;
    }
    dr-election-on-p2p;
    export;
    graceful-restart {
        disable;
        restart-duration seconds;
    }
    import [ policy-names ];
    interface interface-name {
        ... interface-configuration ...
    }
    join-load-balance;
    join-prune-timeout;
    nonstop-routing {
        disable;
    }
    rib-group group-name;
    rp {
    ... rp-configuration ...
    }
    spt-threshold {
        infinity [ policy-names ];
```

```
}
traceoptions {
   file filename <files number> <size size> <world-readable | no-world-readable>;
   flag flag <flag-modifier> <disable>;
}
interface interface-name {
   disable;
   bfd-liveness-detection {
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      version (0 | 1 | automatic);
   }
   hello-interval seconds;
   mode (dense | sparse | sparse-dense);
   neighbor-policy [ policy-names ];
   priority number;
   version version;
}
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-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;
   }
   local {
      disable;
      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;
```

```
                priority number;
            }
        }
        rp-register-policy [ policy-names ];
        static {
            address address {
                group-ranges {
                    destination-ip-prefix</prefix-length>;
                }
                version version;
            }
        }
    }
}
```

Session Announcement Protocol/Session Description Protocol (SAP/SDP)

```
sap {
    disable;
    listen address <port port>;
}
```

[edit routing-instances] Hierarchy Level

```
[edit]
routing-instances routing-instance-name {
    instance-type (forwarding | no-forwarding | virtual-router | vpls | vrf);
    interface interface-name;
    protocols {
        mvpn {
            autodiscovery-only {
                intra-as {
                    inclusive;
                }
            }
            receiver-site;
            sender-site;
            route-target {
                export-target {
                    target target-community;
                    unicast;
                }
                import-target {
                    target {
                        target-value;
                        receiver target-value;
                        sender target-value;
                    }
                    unicast {
                        receiver;
                        sender;
                    }
                }
            }
            unicast-umh-election;
```

```
    }
    pim {
        mdt {
            data-mdt-reuse;
            group-range multicast-prefix;
            threshold {
                group group-address {
                    source source-address {
                        rate threshold-rate;
                    }
                }
            }
            tunnel-limit limit;
        }
        mvpn {
            autodiscovery {
                inet-mdt;
            }
        }
        vpn-group-address class-D-address; /* only for the vrf instance type */
    }
}
provider-tunnel {
    mdt {
        group-range multicast-prefix;
        threshold {
            group group-address {
                source source-address {
                    rate threshold-rate;
                }
            }
            tunnel-limit limit;
        }
    }
    pim-asm {
        group-address address;
    }
    pim-ssm {
        group-address address;
    }
    selective {
        group multicast–prefix</prefix-length {
            source ip–prefix</prefix-length {
                rsvp-te {
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                    static-lsp point-to-multipoint-lsp-name;
                }
                threshold-rate kbps;
            }
            wildcard-source {
                rsvp-te {
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
```

```
                static-lsp point-to-multipoint-lsp-name;
            }
            threshold-rate kbps;
        }
    }
    tunnel-limit number;
    wildcard-group-inet {
        wildcard-source {
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
                static-lsp lsp-name;
            }
            threshold-rate number;
        }
    }
    wildcard-group-inet6 {
        wildcard-source {
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
                static-lsp lsp-name;
            }
            threshold-rate number;
        }
    }
  }
 }
}
```

NOTE: All other statements at the [edit protocols pim] and [edit protocols msdp] hierarchy levels can be configured at the [edit routing-instances *routing-instance-name* protocols pim] and [edit routing-instances *routing-instance-name* protocols msdp] hierarchy levels. The listed statements are valid in a routing instance only.

[edit routing-options] Hierarchy Level

```
[edit]
routing-options {
    multicast {
        asm-override-ssm;
        backup-pe-group group-name {
            backups [ addresses ];
            local-address address;
        }
        flow-map flow-map-name {
            bandwidth (bps | adaptive);
            forwarding-cache {
                timeout (never | minutes);
            }
```

```
            policy [ policy-names ];
            redundant-sources [ addresses ];
        }
        forwarding-cache {
            threshold suppress value <reuse value>;
            timeout minutes;
        }
        interface interface-name {
            maximum-bandwidth bps;
            reverse-oif-mapping;
            subscriber-leave-timer seconds;
        }
        rpf-check-policy [ policy-names ];
        scope scope-name {
            interface [ interface-names ];
            prefix destination-prefix;
        }
        scope-policy [ policy-names ];
        ssm-groups [ addresses ];
        ssm-map ssm-map-name {
            policy [ policy-names ];
            source [ addresses ];
        }
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <disable>;
        }
    }
}
```

[edit logical-systems protocols] Hierarchy Level

```
[edit logical-systems logical-system-name]
protocols {
```

DVMRP

```
dvmrp {
    disable;
    export [ policy-names ];
    import [ policy-names ];
    interface interface-name {
        disable;
        hold-time seconds;
        metric metric;
        mode (forwarding | unicast-routing);
    }
    rib-group group-name;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

IGMP

```
igmp {
    accounting;
    interface interface-name {
```

```
            disable;
            (accounting | no-accounting);
            group-policy [ policy-names ];
            immediate-leave;
            promiscuous-mode;
            ssm-map ssm-map-name;
            static {
                group multicast-group-address {
                    exclude;
                    source ip-address;
                }
            }
            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>;
        }
    }
```

MLD

```
    mld {
        accounting;
        interface interface-name {
            disable;
            (accounting | no-accounting);
            group-policy;
            immediate-leave;
            ssm-map ssm-map-name;
            static {
                group multicast-group-address {
                    exclude;
                    source ip-address;
                }
            }
            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>;
        }
    }
```

MSDP

```
    msdp {
        disable;
        active-source-limit {
            maximum number;
            threshold number;
```

```
    }
    data-encapsulation (disable | enable);
    export [ policy-names ];
    group group-name {
        ... group-configuration ...
    }
    import [ policy-names ];
    local-address address;
    peer address {
        ... peer-configuration ...
    }
    rib-group group-name;
    source ip-prefix</prefix-length> {
        active-source-limit {
            maximum number;
            threshold number;
        }
    }
    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>;
        }
    }
}
```

PGM

```
pgm {
   traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier>;
   }
}
```

PIM

```
pim {
   assert-timeout seconds;
   dense-groups {
      addresses;
   }
   disable;
   dr-election-on-p2p;
   export;
   graceful-restart {
      disable;
      restart-duration seconds;
   }
   import [ policy-names ];
   interface interface-name {
      disable;
      bfd-liveness-detection {
         minimum-interval milliseconds;
         minimum-receive-interval milliseconds;
         minimum-transmit-interval milliseconds;
         multiplier number;
         version (0 | 1 | automatic);
      }
      hello-interval seconds;
      mode (dense | sparse | sparse-dense);
      neighbor-policy [ policy-names ];
      priority number;
      version version;
   }
   join-load-balance;
   join-prune-timeout;
   nonstop-routing {
      disable;
   }
   rib-group group-name;
   rp {
      auto-rp {
         (discovery | mapping);
         (mapping-agent-election | no-mapping-agent-election);
      }
      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;
            }
            local {
               family (inet | inet6) {
                  disable;
                  address address;
                  anycast-pim {
                     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 {
                  version version;
                  group-ranges {
                     destination-ip-prefix</prefix-length>;
                  }
                  traceoptions {
                     file filename <files number> <size size> <world-readable |
                        no-world-readable>;
                     flag flag <flag-modifier> <disable>;
                  }
               }
            }
         }
      }
```

Session Announcement Protocol/Session Description Protocol (SAP/SDP)

```
      sap {
         disable;
         listen address <port port>;
      }
```

[edit logical-systems routing-instances] Hierarchy Level

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

```
routing-instances routing-instance-name {
   interface interface-name;
   instance-type (forwarding | no-forwarding | virtual-router | vpls | vrf);
   protocols {
      pim {
         mdt {
            data-mdt-reuse;
            group-range multicast-prefix;
            threshold {
               group group-address {
                  source source-address {
                     rate threshold-rate;
                  }
               }
            }
            tunnel-limit limit;
         }
         mvpn {
            autodiscovery {
               inet-mdt;
            }
         }
         vpn-group-address class-D-address; /* needed only for vrf instance type */
      }
   }
   provider-tunnel {
      mdt {
         group-range multicast-prefix;
         threshold {
            group group-address {
               source source-address {
                  rate threshold-rate;
               }
            }
            tunnel-limit limit;
         }
      }
      pim-asm {
         group-address address;
      }
      pim-ssm {
         group-address address;
      }
      selective {
         group multicast–prefix</prefix-length {
            source ip–prefix</prefix-length {
               rsvp-te {
                  label-switched-path-template {
                     (default-template | lsp-template-name);
                  }
                  static-lsp point-to-multipoint-lsp-name;
               }
               threshold-rate kbps;
            }
            wildcard-source {
               rsvp-te {
```

```
                label-switched-path-template {
                   (default-template | lsp-template-name);
                }
                static-lsp point-to-multipoint-lsp-name;
             }
             threshold-rate kbps;
          }
       }
       tunnel-limit number;
       wildcard-group-inet {
          wildcard-source {
             rsvp-te {
                label-switched-path-template {
                   (default-template | lsp-template-name);
                }
                static-lsp lsp-name;
             }
             threshold-rate number;
          }
       }
       wildcard-group-inet6 {
          wildcard-source {
             rsvp-te {
                label-switched-path-template {
                   (default-template | lsp-template-name);
                }
                static-lsp lsp-name;
             }
             threshold-rate number;
          }
       }
    }
 }
}
}
}
```

NOTE: Almost all other statements at the [edit protocols pim] and [edit protocols msdp] hierarchy levels can be configured at the [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim] and [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp] hierarchy level. The statements in this section are valid in a routing instance only.

Related Topics

- PIM Mode-Independent Configuration Overview on page 369
- PIM Configuration Statements on page 370
- Changing the PIM Version on page 372
- Configuring Interface Priority to Become the PIM Designated Router on page 372
- Configuring PIM Designated Router Election on Point-to-Point Links on page 373
- Modifying the PIM Hello Interval on page 373
- Configuring Interface-Level PIM Neighbor Policies on page 373
- Configuring a PIM RPF Routing Table on page 374

- Filtering Outgoing PIM Join Messages on page 376
- Filtering Incoming PIM Join Messages on page 377
- Preserving Multicast Performance by Disabling Response to the ping Utility on page 378
- Configuring the BFD Protocol for PIM on page 378
- Overview of BFD Authentication for PIM on page 379
- Configuring BFD Authentication for PIM on page 381
- Configuring Nonstop Active Routing with PIM on page 384
- Disabling PIM on page 385
- Configuring PIM Trace Options on page 388
- Configuring PIM Dense Mode Properties on page 391
- Configuring PIM Sparse-Dense Mode Properties on page 392
- Minimum PIM Sparse Mode Configuration on page 393
- Enabling PIM Sparse Mode on page 394
- Configuring PIM Sparse Mode Graceful Restart on page 395
- Configuring Local PIM RPs on page 396
- Configuring Static PIM RPs on page 399
- Configuring PIM Bootstrap Properties on page 400
- Configuring PIM Auto-RP on page 402
- Configuring Register Message Filtering on a PIM RP or DR on page 407
- Configuring PIM Join Load Balancing on page 408
- Modifying the IGMP Join State Timeout on page 412
- Configuring PIM Embedded RP for IPv6 on page 412
- Configuring the PIM Assert Timeout on page 413
- Configuring the PIM SPT Threshold Policy on page 413
- Configuring Any-Source Multicast for Draft-Rosen VPNs on page 415

[edit logical-systems routing-options] Hierarchy Level

```
[edit logical-systems routing-options]
routing-options {
   multicast {
      asm-override-ssm;
      backup-pe-group group-name {
         backups [ addresses ];
         local-address address;
      }
      flow-map flow-map-name {
         bandwidth (bits-per-second | adaptive);
```

```
            forwarding-cache {
                timeout (never | minutes);
            }
        }
        policy policy-name;
        forwarding-cache {
            threshold suppress value <reuse value>;
            timeout minutes;
        }
        interface interface-name {
            maximum-bandwidth {
                bits-per-second;
            }
        }
        rpf-check-policy [ policy-names ];
        scope scope-name {
            interface [ interface-names ];
            prefix destination-prefix;
        }
        scope-policy policy-name;
        ssm-groups {
            address;
        }
        ssm-map ssm-map-name {
            policy policy-name;
            source addresses;
        }
    }
}
```

[edit bridge-domains bridge-domain-name protocols] Hierarchy Level

Multicast Snooping

```
multicast-snooping-options {
    flood-groups [ ip-addresses ];
    forwarding-cache {
        threshold suppress value <reuse value>;
    }
    graceful-restart <restart-duration seconds>;
    ignore-stp-topology-change;
}
```

IGMP Snooping

```
igmp-snooping {
    immediate-leave;
    interface interface-name {
        group-limit limit;
        host-only-interface;
        immediate-leave;
        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;
    vlan vlan-id {
        immediate-leave;
        interface interface-name {
            group-limit limit;
            host-only-interface;
            immediate-leave;
            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;
    }
}
```

NOTE: All other statements at the [edit bridge-domains *bridge-domain-name* protocols] hierarchy level can be configured at the [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols] hierarchy level. The listed statements are valid on an MX Series router only.

Part 3
IGMP

- IGMP Overview on page 53
- IGMP Configuration Guidelines on page 55
- Summary of IGMP Configuration Statements on page 75

Chapter 5

IGMP Overview

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

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

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

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

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

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

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

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

For information about supported standards for IGMP, see “JUNOS Software Multicast Standards” on page 17.

Chapter 6

IGMP Configuration Guidelines

This chapter describes the following tasks for configuring IGMP:

- Configuring IGMP on page 55
- Minimum IGMP Configuration on page 56
- Enabling IGMP on page 57
- Modifying the IGMP Host-Query Message Interval on page 57
- Modifying the IGMP Query Response Interval on page 57
- Specifying Immediate-Leave Host Removal for IGMP on page 58
- Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 59
- Accepting IGMP Messages from Remote Subnetworks on page 60
- Modifying the IGMP Last-Member Query Interval on page 60
- Modifying the IGMP Robustness Variable on page 61
- Limiting the Maximum IGMP Message Rate on page 62
- Changing the IGMP Version on page 62
- Enabling IGMP Static Group Membership on page 63
- Recording IGMP Join and Leave Events on page 69
- Tracing IGMP Protocol Traffic on page 71
- Disabling IGMP on page 73
- IGMP and Nonstop Active Routing on page 73

Configuring IGMP

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]

For an overview of logical systems and a detailed example of logical system configuration, see the logical systems chapter of the *JUNOS Feature Guide*.

By default, IGMP is automatically enabled on all interfaces on which you configure Protocol Independent Multicast (PIM) and 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 PIM sparse mode, to distribute IGMP group memberships into the multicast routing domain.

Minimum IGMP Configuration

IGMP is automatically enabled on all interfaces on which you configure PIM and all broadcast interfaces when you configure DVMRP. All IGMP configuration statements are optional.

Enabling IGMP

IGMP is automatically enabled on all interfaces on which you configure PIM and all broadcast interfaces when you configure DVMRP.

Optionally, you can specify the interface or interfaces on which to enable IGMP. If you do not specify any interfaces, IGMP is enabled on all interfaces on which you configure PIM and all broadcast interfaces when you configure DVMRP. To enable IGMP explicitly, include the igmp and interface statements:

```
igmp {
    interface interface-name;
}
```

You can include the igmp statement at the following hierarchy levels:

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

For information about specifying interface names, see the *JUNOS Network Interfaces Configuration Guide*.

Modifying the IGMP Host-Query Message Interval

The IGMP querier router periodically sends general host-query messages. These messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.

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

To modify this interval, include the query-interval statement:

```
query-interval seconds;
```

You can include this statement at the following hierarchy levels:

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

The query interval value can be from 1 through 1024 seconds.

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. Varying this interval allows you to adjust the burst peaks of IGMP messages on the subnet.

By default, the query response interval is 10 seconds. To modify this interval, include the query-response-interval statement:

query-response-interval *seconds*;

You can include this statement at the following hierarchy levels:

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

The query response interval can be from 1 through 1024 seconds. It must be less than the host-query message interval.

Specifying Immediate-Leave Host Removal for IGMP

The immediate-leave statement enables you to specify that the router remove a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface.

Use this statement only on IGMP version 2 (IGMPv2) interfaces to which one IGMP host is connected. If more than one IGMP host is connected to a LAN through the same interface, and one host sends a leave group message, the router removes all hosts on the interface from the multicast group. The router loses contact with the hosts that must remain in the multicast group until they send join requests in response to the router's next general group membership query.

When enabled on a router running IGMPv2, and after receiving a leave group membership message from a host associated with the interface, the router immediately removes the group membership from the interface and suppresses the sending of any group-specific queries for the multicast group.

When enabled on a router running IGMP version 3 (IGMPv3), and after the router receives a report containing the type BLOCK_OLD_SOURCES, the router suppresses the sending of group-and-source queries but relies on the JUNOS-supported host tracking mechanism to determine whether or not it removes a particular source group membership from the interface.

To enable immediate leave for an interface, include the immediate-leave statement:

immediate-leave;

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*]

Example: IGMP Immediate Leave

Configure IGMP immediate-leave on the interface where you want hosts removed immediately from multicast groups after the router receives a leave group message:

[edit]

```
protocols {
   igmp {
      interface ge-1/1/1.0 {
         immediate-leave;
      }
   }
}
```

Filtering Unwanted IGMP Reports at the IGMP Interface Level

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. For additional information about how to configure policies, see the *JUNOS Policy Framework Configuration Guide*.

To enable IGMP report filtering for an interface, include the group-policy statement:

```
group-policy [ policy-names ];
```

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*]

Example: IGMP Report Filtering

Configure IGMP report filtering on the interface where you prefer to not receive specific group or (source, group) reports:

```
[edit]
protocols {
   igmp {
      interface ge-1/1/1.0 {
         group-policy reject_policy;
      }
   }
}
```

Configure either an IGMP version 2 or IGMP version 3 policy:

```
policy-options { #IGMPv2 policy
   policy-statement reject_policy {
      from {
         route-filter 224.1.1.1/32 exact;
```

```
        }
        then reject;
    }
    policy-statement reject_policy { #IGMPv3 policy
        from {
            route-filter 224.1.1.1/32 exact;
            source-address-filter 10.1.0.0/16 orlonger;
        }
        then reject;
    }
}
```

Accepting IGMP Messages from Remote Subnetworks

By default, IGMP interfaces accept IGMP messages only from the same subnetwork. Including the promiscuous-mode statement enables the router to accept IGMP messages from different subnetworks.

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.

To enable promiscuous mode for an interface, include the promiscuous-mode statement:

```
promiscuous-mode;
```

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*]

Example: IGMP Promiscuous Mode

Enable IGMP promiscuous mode on the interface you want to accept IGMP reports from indirectly connected subnets:

```
[edit]
protocols {
    igmp {
        interface ge-1/1/1.0 {
            promiscuous-mode;
        }
    }
}
```

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 lower this interval to reduce the amount of time it takes a router to detect the loss of the last member of a group.

The default last-member query interval is 1 second. To modify this interval, include the query-last-member-interval statement:

query-last-member-interval *seconds*;

You can include this statement at the following hierarchy levels:

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

The last-member query interval can be from 1 through 1024 seconds.

Modifying the IGMP Robustness Variable

The IGMP robustness variable provides fine-tuning to allow for expected packet loss on a subnet. The value of the robustness variable is used in calculating the following IGMP message intervals:

- Group member interval—Amount of time that must pass before a multicast router decides 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 decides 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. You might want to increase this value if you expect a subnet to lose packets. To change the value of the robustness variable, include the robust-count statement:

robust-count *number*;

You can include this statement at the following hierarchy levels:

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

The number can be from 2 through 10.

Limiting the Maximum IGMP Message Rate

This section describes how to change the limit for the maximum number of IGMP packets transmitted in one 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 Topics

- maximum-transmit-rate

Changing the IGMP Version

By default, the router runs IGMPv2. Routers running different versions of IGMP negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version.

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

NOTE: 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.

To change to IGMPv3 (for source-specific multicast [SSM] functionality), include the version statement:

```
version 3;
```

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*]

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

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

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.

In this example, you create static group 225.1.1.1.

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

   user@host# **set protocols igmp 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 routers 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 needing 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.

   user@host# **set protocols igmp interface fe-0/1/2 static group 225.1.1.1 group-count 3**

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

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

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

3. After you have committed the configuration and 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 needing to configure each receiver separately and you do not want the host addresses to be sequential.

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

1. 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.

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

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

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

   ```
   interface fe-0/1/2.0 {
      static {
         group 225.1.1.1 {
            group-increment 0.0.0.2;
            group-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 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 configure the multicast source address to 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. Configure the source address by including the source statement and specifying the IPv4 address of the source host.

   user@host# **set protocols igmp 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 {
      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. 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.

user@host# **set protocols igmp 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 {
   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: 225.1.1.3
         Source: 10.0.0.3
         Last reported by: Local
         Timeout: 0 Type: Static
     Group: 225.1.1.5
         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 needing to configure each receiver separately and you do not want the host 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.

user@host# **set protocols igmp 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 {
   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 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: 225.1.1.5
         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. Configure a multicast static group to operate in exclude mode by including the exclude statement and specifying the IPv4 source address that should be excluded.

```
user@host# set protocols igmp 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 {
   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
```

Similar configuration is available for IPv6 multicast traffic using the MLD protocol.

Related Topics

- group
- group-count
- group-increment
- static
- source-count
- source-increment
- Enabling MLD Static Group Membership on page 105

Recording IGMP Join and Leave Events

You can configure the router to record IGMP join and leave events and disable event recording for individual interfaces using the accounting statement.

Enabling IGMP Accounting on the Entire Routing System

To enable the recording of various IGMP join and leave events on the entire routing system, include the accounting statement:

```
accounting;
```

You can include this statement at the following hierarchy levels:

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

Table 5 on page 70 describes the recordable IGMP join and leave events and IGMP accounting enabled or disabled events for an IGMP interface:

Table 5: IGMP Event Messages

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

After enabling IGMP accounting, you must configure the router to filter the recorded information to a file or display it to a terminal. Optionally, you can archive the events file.

For detailed information about defining and archiving system logs, see the *JUNOS System Basics Configuration Guide*.

Enabling or Disabling IGMP Accounting on Individual Interfaces

To enable or disable the IGMP join and leave event recording for individual interfaces, include the accounting or no-accounting statement:

```
(accounting | no-accounting);
```

You can include these statements at the following hierarchy levels:

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

Example: Recording and Archiving IGMP Join and Leave Events

Enable IGMP accounting on the router and disable event recording on an individual interface:

```
[edit]
protocols {
   igmp {
      accounting; # accounting enabled on the IGMP router
      interface ge-1/1/1.0 {
         no-accounting; # accounting disabled on interface ge-1/1/1.0
      }
   }
}
```

Enable IGMP accounting on an individual interface only:

```
[edit]
protocols {
   igmp {
      interface ge-1/2/1.0 {
         accounting; # accounting enabled on interface ge-1/2/1.0
      }
   }
}
```

Filter the events to a system log file and periodically archive the file:

```
[edit system]
syslog {
   file igmp-events { # define the system log filename
      any info;
      match ".*RPD_IGMP_JOIN.* | .*RPD_IGMP_LEAVE.* | .*RPD_IGMP_ACCOUNTING.*
         | .*RPD_IGMP_MEMBERSHIP_TIMEOUT.*";
         /* define the events you want to record */
      archive {
         size 10000; # rotate the file size when it reaches 10 KB
         files 3; # keep the last three files
         archive-sites {
            "ftp://regress@mars//var/tmp" password "anonymous";
            "ftp://regress@saturn//var/tmp" password "test";
            }
            transfer-interval 24 h; #upload the file every 24 hours
         start-time 2009–01–07:12:30; # date and time at which the upload occurs
      }
   }
}
```

Tracing IGMP Protocol Traffic

To trace IGMP protocol traffic, specify options to the **traceoptions** statement at the [edit routing-options] or [edit logical-systems *logical-system-name* routing-options]

hierarchy level. Options applied at the routing options level trace all packets, and options applied at the protocol level trace only IGMP traffic.

You can specify IGMP-specific options by including the traceoptions statement:

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

You can specify the following IGMP-specific options in the IGMP flag statement:

- leave—Trace leave-group messages (for IGMPv2 only).
- mtrace—Trace mtrace packets. Use the mtrace command to troubleshoot the software.
- packets—Trace all IGMP packets.
- query—Trace IGMP membership query messages, including general and group-specific queries.
- report—Trace membership report messages.

To trace the paths of multicast packets, use the mtrace command, as described in the *JUNOS System Basics and Services Command Reference*.

For information about tracing and global tracing options, see the *JUNOS Routing Protocols Configuration Guide*.

Example: Tracing IGMP Protocol Traffic

Trace only unusual or abnormal operations to the file routing-log, and trace all IGMP packets to the file igmp-log:

```
[edit]
routing-options {
   traceoptions {
      file routing-log;
      flag errors;
   }
}
protocols {
   igmp {
      traceoptions {
         file igmp-log;
         flag packets;
      }
   }
}
```

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*]

IGMP and Nonstop Active Routing

Nonstop active routing configurations include two Routing Engines that share information so that routing is not interrupted during Routing Engine failover. These NSR configurations include passive support with IGMP in connection with PIM. The master Routing Engine uses IGMP to determine its PIM multicast state, and this IGMP-derived information is replicated on the backup Routing Engine. IGMP on the new master Routing Engine (after-failover) relearns the state information quickly through IGMP operation. In the interim, the new master Routing Engine retains the IGMP-derived PIM state as received by the replication process from the old master Routing Engine. This state information times out unless refreshed by IGMP on the new master Routing Engine. No additional IGMP configuration is required.

For more information about IGMP and nonstop active routing. see “Configuring Nonstop Active Routing with PIM” on page 384 and the *JUNOS High Availability Configuration Guide*.

Chapter 7

Summary of IGMP Configuration Statements

The following sections explain each of the Internet Group Management Protocol (IGMP) configuration statements. The statements are organized alphabetically.

accounting (Per Interface)

Syntax (accounting | no-accounting);

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

Release Information Statement introduced in JUNOS Release 8.5.

Description Enable or disable the collection of IGMP join and leave event statistics for an interface.

Usage Guidelines See "Recording IGMP Join and Leave Events" on page 69.

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

accounting (Protocol)

Syntax accounting;

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

Release Information Statement introduced in JUNOS Release 8.5.

Description Enable the collection of IGMP join and leave event statistics on the system.

Usage Guidelines See "Recording IGMP Join and Leave Events" on page 69.

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

disable

Syntax disable;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Disable IGMP on the system.

Usage Guidelines See "Disabling IGMP" on page 73.

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

exclude

Syntax exclude;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in JUNOS 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 Topics Enabling IGMP Static Group Membership on page 63

group

Syntax

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

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

Release Information Statement introduced before JUNOS Release 7.4.

Description IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.

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

The remaining statements are explained separately.

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

Related Topics Enabling IGMP Static Group Membership on page 63

group-count

Syntax group-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in JUNOS Release 9.6.

Description Specify the number of static groups to be created.

Options *number*—Number of static groups.
Default:
Range: 1 through 512

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

Related Topics Enabling IGMP Static Group Membership on page 63

group-increment

Syntax group-increment *increment*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in JUNOS Release 9.6.

Description Configure the number of times the address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.

Options *increment*—Number of times the address should be incremented.
Default: 0.0.0.1
Range: 0.0.0.1 through 255.255.255.255

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

Related Topics Enabling IGMP Static Group Membership on page 63

group-policy

Syntax group-policy [*policy-names*];

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

Release Information Statement introduced in JUNOS Release 9.1.

Description When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report).

Usage Guidelines See “Filtering Unwanted IGMP Reports at the IGMP Interface Level” on page 59.

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

igmp

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

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).

Usage Guidelines

See “Enabling IGMP” on page 57.

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

immediate-leave

Syntax immediate-leave;

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

Release Information Statement introduced in JUNOS Release 8.3.

Description When this statement is enabled on a router running IGMP version 2 (IGMPv2), after the router receives a leave group membership message from a host associated with the interface, the router immediately removes the group membership from the interface and suppresses the sending of any group-specific queries for the multicast group.

When this statement is enabled on a router running IGMP version 3 (IGMPv3), after the router receives a report with the type BLOCK_OLD_SOURCES, the router suppresses the sending of group-and-source queries but relies on the JUNOS-supported host tracking mechanism to determine whether or not it removes a particular source group membership from the interface.

NOTE: When issuing this command on IGMPv2 interfaces, ensure that the IGMP interface has only one IGMP host connected. If more than one IGMPv2 host is connected to a LAN through the same interface, and one host sends a done message, the router removes all hosts on the interface from the multicast group. The router loses contact with the hosts that properly remain in the multicast group until they send join requests in response to the next general multicast listener query from the router.

Usage Guidelines See “Specifying Immediate-Leave Host Removal for IGMP” on page 58.

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

interface

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Enable IGMP on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**. For details about specifying interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

The remaining statements are explained separately.

Usage Guidelines

See "Enabling IGMP" on page 57.

Required Privilege Level

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

maximum-transmit-rate

Syntax maximum-transmit-rate *packets-per-second*;

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

Release Information Statement introduced in JUNOS Release 9.3.

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 Topics
- Limiting the Maximum IGMP Message Rate on page 62

no-accounting

See accounting.

oif-map

Syntax oif-map *map-name*;

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

Release Information Statement introduced in JUNOS Release 9.6.

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 Topics

- Configuring Multicast Outgoing Interface Mapping on page 220

passive (IGMP)

Syntax passive <allow-receive> <send-general-query> <send-group-query>;

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

Release Information Statement introduced in JUNOS Release 9.6.
allow-receive, send-general-query, and send-group-query options were added in JUNOS Release 10.0.

Description Specify that IGMP run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as IGMP reports, queries, and leaves.

NOTE: You can selectively activate up to two out of the three available options for the passive statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the passive statement.

Options allow-receive—Enables IGMP to receive control traffic on the interface.

send-general-query—Enables IGMP to send general queries on the interface.

send-group-query—Enables IGMP to send group-specific and group-source-specific queries on the interface.

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

Related Topics

- Configuring Multicast Outgoing Interface Mapping on page 220
- For general information about configuring IGMP, see the *JUNOS Multicast Protocols Configuration Guide*.

promiscuous-mode

Syntax promiscuous-mode;

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

Release Information Statement introduced in JUNOS Release 8.3.

Description Specify that the interface accepts IGMP reports from hosts on any subnetwork.

Usage Guidelines See "Accepting IGMP Messages from Remote Subnetworks" on page 60.

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

query-interval

Syntax query-interval *seconds*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description How often the querier router sends general host-query messages.

Options *seconds*—Time interval.
Range: 1 through 1024
Default: 125 seconds

Usage Guidelines See "Modifying the IGMP Host-Query Message Interval" on page 57.

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

Related Topics
- query-last-member-interval
- query-response-interval

query-last-member-interval

Syntax query-last-member-interval *seconds*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description How often the querier router 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 1024
Default: 1 second

Usage Guidelines See "Modifying the IGMP Last-Member Query Interval" on page 60.

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

Related Topics
- query-interval
- query-response-interval

query-response-interval

Syntax query-response-interval *seconds*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description How long the querier router waits to receive a response to a host-query message from a host.

Options *seconds*—Time interval. This interval must be less than the interval between general host-query messages.
Range: 1 through 1024
Default: 10 seconds

Usage Guidelines See "Modifying the IGMP Query Response Interval" on page 57.

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

Related Topics
- query-last-member-interval
- query-interval

robust-count

Syntax robust-count *number*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Dimensionless factor used to calculate several IGMP message intervals. Raised for higher expected packet loss on a subnet.

Options *number*—Robustness variable.
Range: 2 through 10
Default: 2

Usage Guidelines See "Modifying the IGMP Robustness Variable" on page 61.

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

source

Syntax
```
source ip-address {
   source-count number;
   source-increment increment;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description IP version 4 (IPv4) unicast source address for the multicast group being statically configured on an interface.

Options *ip-address*—IPv4 unicast address.

The remaining statements are explained separately.

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

Related Topics Enabling IGMP Static Group Membership on page 63

source-count

Syntax source-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address* source],
[edit protocols igmp interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in JUNOS Release 9.6.

Description Configure the number of multicast source addresses that should be accepted for each static group created.

Options *number*—Number of source addresses.
Default: 1
Range: 1 through 1024

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

Related Topics Enabling IGMP Static Group Membership on page 63

source-increment

Syntax source-increment *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address* source],
[edit protocols igmp interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in JUNOS Release 9.6.

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 Topics Enabling IGMP Static Group Membership on page 63

ssm-map

Syntax ssm-map *ssm-map-name*;

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

Release Information Statement introduced in JUNOS Release 7.4.

Description Apply an SSM map to an IGMP interface.

Options *ssm-map-name*—Name of SSM map.

Usage Guidelines See "Example: Configuring SSM Mapping" on page 208.

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

static

Syntax

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

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

Release Information Statement introduced before JUNOS Release 7.4.

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

The remaining statements are explained separately.

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

Related Topics Enabling IGMP Static Group Membership on page 63

traceoptions

Syntax

```
traceoptions {
   file filename <files number> <size size> <world-readable | no-world-readable>;
   flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Configure IGMP tracing options.

To specify more than one tracing operation, include multiple flag statements.

To trace the paths of multicast packets, use the mtrace command, as described in the *JUNOS System Basics and Services Command Reference*.

Default

The default IGMP trace options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place tracing output in the file igmp-log.

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.
Range: 2 through 1000 files
Default: 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.

IGMP Tracing Flags

- leave—Leave group messages (for IGMP version 2 only).
- mtrace—Mtrace packets. Use the mtrace command to troubleshoot the software.
- packets—All IGMP packets.

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

Global Tracing Flags

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

  Default: If you do not specify this option, only unusual or abnormal operations are traced.
- policy—Policy operations and actions
- route—Routing table changes
- state—State transitions
- task—Interface transactions and processing
- timer—Timer usage

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

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

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.
Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

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

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

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

If you specify a maximum file size, you must also include the files statement to specify the maximum number of trace files.
Syntax: *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB
Range: 10 KB through the maximum file size supported on your system
Default: 1 MB

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

Usage Guidelines See "Tracing IGMP Protocol Traffic" on page 71.

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.

version

Syntax version *version*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the version of IGMP.

Options *version*—IGMP version number.
Range: 1, 2, or 3
Default: IGMP version 2

Usage Guidelines See "Changing the IGMP Version" on page 62.

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

Part 4
MLD

- MLD Overview Introduction on page 95
- MLD Configuration Guidelines on page 99
- Summary of MLD Configuration Statements on page 117

Chapter 8

MLD Overview Introduction

MLD Overview

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 router maintains a list of host multicast addresses that have listeners for each subnet, as well as a timer for each address. However, the router does not need to know the address of the listeners—just the address of the hosts. The router provides addresses to the multicast routing protocol it uses; this ensures that multicast packets are delivered to all subnets 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 routers and hosts that need to receive IP multicast traffic. The JUNOS Software 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 from which it is interested in receiving the multicast group traffic. Exclude mode works the opposite of include mode. It allows the receiver to specify the sources or sources from which it is not interested in receiving the multicast group traffic from.

For information about supported standards for MLD, see “IP Multicast Specifications” on page 18.

For each attached network, a multicast router can be either a querier or a nonquerier. A querier router, usually one per subnet, solicits group membership information by transmitting MLD queries. When a host reports to the querier router that it has interested listeners, the querier router forwards the membership information to the rendezvous point (RP) router 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 router. The RPT is the initial path used by the sender to transmit information to the interested listeners. For more information about PIM distribution trees, see “PIM Sparse Mode” on page 342. Nonquerier routers do not transmit MLD queries on a subnet but can do so if the querier router goes down.

All MLD-configured routers start up as querier routers on each attached subnet (see Figure 4 on page 96). The querier router on the right is the receiver's DR.

Figure 4: Routers Start Up on a Subnet

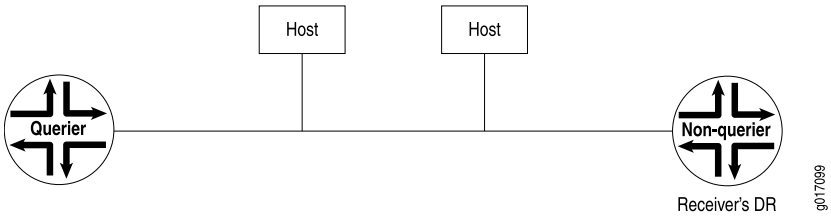

To elect the querier router, the routers exchange query messages containing their IPv6 source addresses. If a router hears a query message whose IPv6 source address is numerically lower than its own selected address, it becomes a nonquerier. In Figure 5 on page 96, the router on the left has a source address numerically lower than the one on the right and therefore becomes the querier router.

NOTE: In the practical application of MLD, several routers on a subnet are nonqueriers. If the elected querier router goes down, query messages are exchanged among the remaining routers. The router with the lowest IPv6 source address then becomes the new querier router. Note that 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 5: Querier Router Is Determined

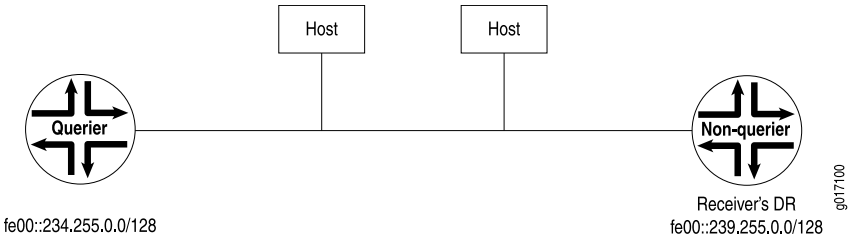

The querier router sends general MLD queries on the link-scope all-nodes multicast address FF02::1 at short intervals to all attached subnets to solicit group membership information (see Figure 6 on page 96). Within the query message is the *maximum response delay* value, specifying the maximum allowed delay for the host to respond with a report message.

Figure 6: General Query Message Is Issued

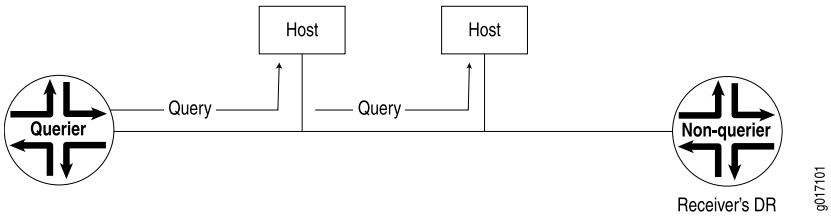

If interested listeners are attached to the host receiving the query, the host sends a report containing the host's IPv6 address to the router (see Figure 7 on page 97). If the reported address is not yet in the router's list of multicast addresses with interested listeners, the address is added to the 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 7: Reports Are Received by the Querier Router

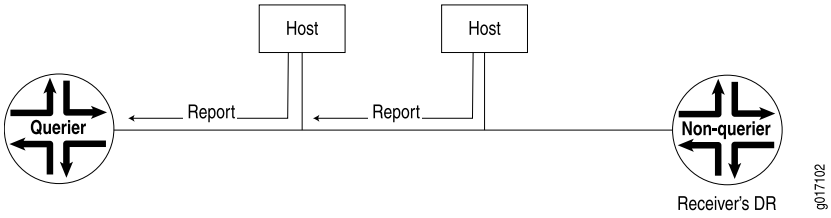

If the host has no interested multicast listeners, it sends a done message to the querier router. On receipt, the querier router issues a multicast address–specific query containing the last listener query interval value to the multicast address of the host. If the router does not receive a report from the multicast address, it removes the multicast address from the list and notifies the RP in the PIM domain of its removal (see Figure 8 on page 97).

Figure 8: Host Has No Interested Receivers and Sends a Done Message to Router

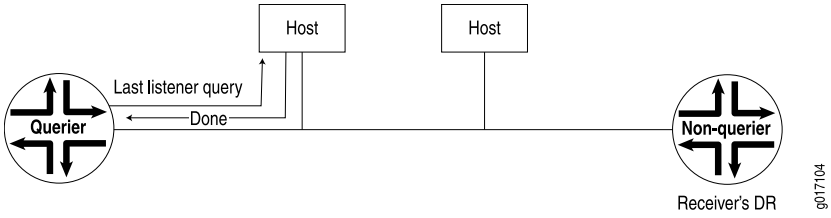

If a done message is not received by the querier router, the querier router continues to send multicast address–specific queries. If the timer set for the address on receipt of the last report expires, the querier router assumes there are no longer interested listeners present on that subnet, removes the multicast address from the list, and notifies the RP in the PIM domain of its removal (see Figure 9 on page 97).

Figure 9: Host Address Timer Expires and Address Is Removed from Multicast Address List

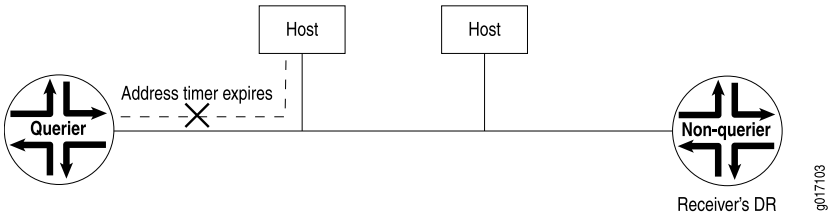

Chapter 9

MLD Configuration Guidelines

This chapter describes the following tasks for configuring MLD:

- Configuring MLD on page 99
- Minimum MLD Configuration on page 100
- Enabling MLD on page 100
- Modifying the MLD Version on page 101
- Modifying the MLD Host-Query Message Interval on page 101
- Modifying the MLD Query Response Interval on page 101
- Modifying the MLD Last-Member Query Interval on page 102
- Specifying Immediate-Leave Host Removal for MLD on page 102
- Filtering Unwanted MLD Reports at the MLD Interface Level on page 103
- Modifying the MLD Robustness Variable on page 104
- Limiting the Maximum MLD Message Rate on page 105
- Enabling MLD Static Group Membership on page 105
- Recording MLD Join and Leave Events on page 112
- Tracing MLD Protocol Traffic on page 114
- Disabling MLD on page 115

Configuring MLD

To configure the Multicast Listener Discovery (MLD) protocol, include the mld statement:

```
mld {
    accounting;
    interface interface-name {
        disable;
        (accounting | no-accounting);
        group-policy;
        immediate-leave;
        ssm-map ssm-map-name;
        static {
            group multicast-group-address {
                exclude;
                source ip-address;
```

```
            }
        }
        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]

For an overview of logical systems and a detailed example of logical system configuration, see the logical systems chapter of the *JUNOS Feature Guide*.

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

Minimum MLD Configuration

MLD is automatically enabled on all broadcast interfaces when you configure PIM or DVMRP. All MLD configuration statements are optional.

Enabling MLD

MLD is automatically enabled on all broadcast interfaces when you configure PIM or DVMRP.

To enable MLD explicitly, include the mld statement. Optionally, you can specify the interface or interfaces on which to enable MLD. If you do not specify any interfaces, MLD is enabled on all interfaces.

```
mld {
    interface interface-name;
}
```

You can include this statement at the following hierarchy levels:

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

For information about specifying interface names, see the sections about interface naming in the *JUNOS Network Interfaces Configuration Guide*.

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.

```
version 2;
```

You can include this statement at the following hierarchy levels:

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

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

Related Topics

- Source-Specific Multicast Groups Overview on page 203
- Configuring Source-Specific Multicast Groups on page 204
- Example: Configuring an SSM-Only Domain on page 206
- Example: Configuring PIM SSM on a Network on page 206
- Example: Configuring SSM Mapping on page 208

Modifying the MLD Host-Query Message Interval

The MLD querier router periodically sends general host-query messages. These messages solicit group membership information and are sent to the link-scope all-nodes address FF02::1.

By default, host-query messages are sent every 125 seconds. You can change the number of MLD messages sent on the subnet by changing the query interval value. The larger the value, the less often MLD queries are sent.

To modify the interval value, include the query-interval statement:

```
query-interval seconds;
```

You can include this statement at the following hierarchy levels:

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

The query interval value can range from 1 through 1024 seconds.

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. Varying this interval allows you to adjust the burst peaks of MLD messages on the subnet. Larger intervals create more widely spaced node responses and result in less bursty traffic.

By default, the query response interval is 10 seconds. To modify this interval, include the query-response-interval statement:

query-response-interval *seconds*;

You can include this statement at the following hierarchy levels:

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

The query response interval can range from 1 through 1024 seconds. It must be less than the host-query message interval.

Modifying the MLD 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 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.

The default last-member query interval is 1 second. To modify this interval, include the query-last-member-interval statement:

query-last-member-interval *seconds*;

You can include this statement at the following hierarchy levels:

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

The last-member query interval can range from 1 through 1024 seconds.

Specifying Immediate-Leave Host Removal for MLD

The immediate-leave statement enables you to specify that the router remove a host from the multicast group as soon as the router receives a multicast leave group message from a host associated with this interface.

Use this statement only on MLD interfaces to which one MLD host is connected. If more than one MLD host is connected to a LAN through the same interface, and one host sends a done message, the router removes all hosts on the interface from the multicast group. The router loses contact with the hosts that properly remain in the multicast group until they send join requests in response to the router's next general multicast listener query.

When this statement is enabled on a router running MLDv1, after the router receives a multicast listener done message from a host associated with the interface, the

router immediately removes the group membership from the interface and suppresses the sending of any group-specific queries for the multicast group.

When this statement is enabled on a router running MLDv2, after the router receives a report with the type BLOCK_OLD_SOURCES, the router suppresses the sending of group-and-source queries but relies on the JUNOS-supported host tracking mechanism to determine whether or not it removes a particular source group membership from the interface.

To enable immediate leave for an interface, include the immediate-leave statement:

```
immediate-leave;
```

You can include this statement at the following hierarchy levels:

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

Example: MLD Immediate Leave

Configure MLD immediate leave on the interface where you want hosts removed immediately from multicast groups after the router receives a leave group message:

```
[edit ]
protocols {
   mld {
      interface fe-1/0/1.0 {
         immediate-leave;
      }
   }
}
```

Filtering Unwanted MLD Reports at the MLD Interface Level

The group-policy statement enables you to filter unwanted IGMP reports at the interface level. When this statement is enabled on a router running MLD version 1 (MLDv1), 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). When this statement is enabled on a router running MLD version 2 (MLDv2), 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. For additional information about how to configure policies, see the *JUNOS Policy Framework Configuration Guide*.

To enable MLD report filtering for an interface, include the group-policy statement:

```
group-policy [ policy-names ];
```

You can include this statement at the following hierarchy levels:

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

Example: MLD Report Filtering

Configure MLD report filtering on the interface where you prefer not to receive specific group or (source, group) reports:

```
[edit]
protocols {
   mld {
      interface ge-1/1/1.0 {
         group-policy reject_policy;
      }
   }
}
```

Configure either an MLDv1 or MLDv2 policy:

```
policy-options { #MLDv1 policy
   policy-statement reject_policy {
      from {
         route-filter fec0:1:1:4::/64 exact;
      }
      then reject;
   }
   policy-statement reject_policy { #MLDv2 policy
      from {
         route-filter fec0:1:1:4::/64 exact;
         source-address-filter fe80::2e0:81ff:fe05:1a8d/32 orlonger;
      }
      then reject;
   }
}
```

Modifying the MLD Robustness Variable

The MLD robustness variable provides fine-tuning to allow for expected packet loss on a subnet. 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 decides 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 decides 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. You might want to increase this value if you expect a subnet to lose packets. To change the value of the robustness variable, include the robust-count statement:

robust-count *number*;

You can include this statement at the following hierarchy levels:

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

The number can be from 2 through 10.

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.

Related Topics
- maximum-transmit-rate

Enabling MLD Static Group Membership

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.

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 ff02::1:ff05:1a8d.

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

```
user@host# set protocols mld interface fe-0/1/2 static group ff02::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 ff02::1:ff05:1a8d;
   }
}
```

3. After you have committed the configuration and the source is sending traffic, use the show mld group command to verify that static group ff02::1:ff05:1a8 has been created.

```
user@host> show mld group

Interface: fe-0/1/2
Group: ff02::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.

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 needing 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.

```
user@host# set protocols mld interface fe-0/1/2 static group ff02::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 ff02::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 ff02::1:ff05:1a8d, ff02::1:ff05:1a8e, and ff02::1:ff05:1a8f have been created.

```
user@host> show mld group


Interface: fe-0/1/2
     Group: ff02::1:ff05:1a8d
         Source: fe80::2e0:81ff:fe05:1a8d
         Last reported by: Local
         Timeout: 0 Type: Static
Interface: fe-0/1/2
     Group: ff02::1:ff05:1a8e
         Source: fe80::2e0:81ff:fe05:1a8d
         Last reported by: Local
         Timeout: 0 Type: Static
Interface: fe-0/1/2
     Group: ff02::1:ff05:1a8f
         Source: fe80::2e0:81ff:fe05:1a8d
         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 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.

   user@host# **set protocols mld interface fe-0/1/2 static group ff02::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 ff02::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 ff02::1:ff05:1a8d, ff02::1:ff05:1a8f, and ff02::1:ff05:1a91 have been created.

```
user@host> show mld group


Interface: fe-0/1/2
     Group: ff02::1:ff05:1a8d
         Source: fe80::2e0:81ff:fe05:1a8d
         Last reported by: Local
         Timeout: 0 Type: Static
Interface: fe-0/1/2
     Group: ff02::1:ff05:1a8f
         Source: fe80::2e0:81ff:fe05:1a8d
         Last reported by: Local
         Timeout: 0 Type: Static
Interface: fe-0/1/2
     Group: ff02::1:ff05:1a91
         Source: fe80::2e0:81ff:fe05:1a8d
         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 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 ff02::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.

   ```
   user@host# set protocols mld interface fe-0/1/2 static group ff02::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 ff02::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 ff02::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: ff02::1:ff05:1a8d
         Source: fe80::2e0:81ff:fe05:1a8d
         Last reported by: Local
         Timeout: 0 Type: Static
```

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 ff02::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.

   user@host# **set protocols mld interface fe-0/1/2 static group ff02::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 ff02::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 ff02::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: ff02::1:ff05:1a8d
            Source: fe80::2e0:81ff:fe05:1a8d
            Last reported by: Local
            Timeout: 0 Type: Static
   Interface: fe-0/1/2
        Group: ff02::1:ff05:1a8d
            Source: fe80::2e0:81ff:fe05:1a8e
            Last reported by: Local
            Timeout: 0 Type: Static
   Interface: fe-0/1/2
        Group: ff02::1:ff05:1a8d
   ```

```
            Source: fe80::2e0:81ff:fe05:1a8f
            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 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 ff02::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.

   user@host# **set protocols mld interface fe-0/1/2 static group ff02::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 ff02::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 ff02::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: ff02::1:ff05:1a8d
            Source: fe80::2e0:81ff:fe05:1a8d
            Last reported by: Local
            Timeout: 0 Type: Static
   Interface: fe-0/1/2
        Group: ff02::1:ff05:1a8d
            Source: fe80::2e0:81ff:fe05:1a8f
            Last reported by: Local
            Timeout: 0 Type: Static
   Interface: fe-0/1/2
   ```

```
        Group: ff02::1:ff05:1a8d
            Source: fe80::2e0:81ff:fe05:1a91
            Last reported by: Local
            Timeout: 0 Type: Static

    Interface: fe-0/1/2
    Group: ff02::1:ff05:1a8d
    Group mode: Include
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
    Group: ff02::1:ff05:1a8d
    Group mode: Include
    Source: fe80::2e0:81ff:fe05:1a8f
    Last reported by: Local
    Timeout: 0 Type: Static
    Group: ff02::1:ff05:1a8d
    Group mode: Include
    Source: fe80::2e0:81ff:fe05:1a91
    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 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 ff02::1:ff05:1a8d.

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

   user@host# **set protocols mld interface fe-0/1/2 static group ff02::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 ff02::1:ff05:1a8d {
            exclude;
            source fe80::2e0:81ff:fe05:1a8d;
         }
      }
   ```

```
}
```

3. After you have committed the configuration and the source is sending traffic, use the show mld group detail command to verify that static group ff02::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: ff02::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.

Related Topics Enabling IGMP Static Group Membership on page 63

Recording MLD Join and Leave Events

You can use the accounting statement to configure the router to record MLD join and leave events and disable event recording for individual interfaces.

Enabling MLD Accounting on the Entire Routing System

To enable the recording of various MLD join and leave events on the entire routing system, include the accounting statement:

accounting;

You can include this statement at the following hierarchy levels:

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

Table 6 on page 112 describes the recordable MLD join and leave events and MLD accounting enabled or disabled events for an MLD interface:

Table 6: MLD Event Messages

ERRMSG	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 6: MLD Event Messages *(continued)*

ERRMSG	Definition
RPD_MLD_MEMBERSHIP_TIMEOUT	Records MLD membership timeout events.

After enabling MLD accounting, you must configure the router to filter the recorded information to a file or display it to a terminal. Optionally, you can archive the events file.

For detailed information about defining and archiving system logs, see the *JUNOS System Basics Configuration Guide*.

Enabling or Disabling MLD Accounting on Individual Interfaces

To enable or disable the MLD join and leave event recording for individual interfaces, include the accounting or no-accounting statement:

```
(accounting | no-accounting);
```

You can include these statements at the following hierarchy levels:

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

Example: Recording and Archiving MLD Join and Leave Events

Enable MLD accounting on the router and disable event recording on an individual interface:

```
[edit]
protocols {
   mld {
      accounting; # accounting enabled on the MLD router
      interface ge-1/1/1.0 {
         no-accounting; # accounting disabled on interface ge-1/1/1.0
      }
   }
}
```

Enable MLD accounting on an individual interface only:

```
[edit]
protocols {
   mld {
      interface ge-1/2/1.0 {
         accounting; # accounting enabled on interface ge-1/2/1.0
      }
   }
}
```

Filter the events to a system log file and periodically archive the file:

```
[edit system]
syslog {
   file mld-events { # define the system log filename
      any info;
      match ".*RPD_MLD_JOIN.* | .*RPD_MLD_LEAVE.* | .*RPD_MLD_ACCOUNTING.*
         | .*RPD_MLD_MEMBERSHIP_TIMEOUT.*";
         /* define the events you want to record */
      archive {
         size 10000; # rotate the file size when it reaches 10 KB
         files 3; # keep the last three files
         archive-sites {
            "ftp://regress@mars//var/tmp" password "anonymous";
            "ftp://regress@saturn//var/tmp" password "test";
            }
            transfer-interval 24 h; # upload the file every 24 hours
         start-time 2009–01–07:12:30; # date and time at which the upload occurs
      }
   }
}
```

Tracing MLD Protocol Traffic

To trace MLD protocol traffic, you can specify options in the global traceoptions statement at the [edit routing-options] or [edit logical-systems *logical-system-name* routing-options] hierarchy level. Options applied at the routing options level trace all packets, and options applied at the protocol level trace only IGMP traffic. You can specify MLD-specific options by including the traceoptions statement:

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

You can specify the following MLD-specific options in the MLD flag statement:

- leave—Trace leave-group messages (for version 2 only).
- mtrace—Trace mtrace packets. Use the mtrace command to troubleshoot the software.
- packets—Trace all MLD packets.
- query—Trace MLD membership query messages, including general and group-specific queries.
- report—Trace membership report messages.

To trace the paths of multicast packets, use the mtrace command, as described in the *JUNOS System Basics and Services Command Reference*.

For information about tracing and global tracing options, see the *JUNOS Routing Protocols Configuration Guide*.

Example: Tracing MLD Protocol Traffic

Trace only unusual or abnormal operations to the file routing-log, and trace all MLD packets to the file mld-log:

```
[edit]
routing-options {
   traceoptions {
      file routing-log;
      flag errors;
   }
}
protocols {
   mld {
      traceoptions {
         file mld-log;
         flag packets;
      }
   }
}
```

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]

For information about specifying interface names, see the sections about interface naming in the *JUNOS Network Interfaces Configuration Guide*.

Chapter 10

Summary of MLD Configuration Statements

The following sections explain each of the Multicast Listener Discovery (MLD) configuration statements. The statements are organized alphabetically.

accounting (Per Interface)

Syntax (accounting | no-accounting);

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

Release Information Statement introduced in JUNOS Release 9.1.

Description Enable or disable the collection of MLD join and leave event statistics for an interface.

Usage Guidelines See "Recording MLD Join and Leave Events" on page 112.

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

accounting (Protocol)

Syntax accounting;

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

Release Information Statement introduced in JUNOS Release 9.1.

Description Enable the collection of MLD join and leave event statistics on the system.

Usage Guidelines See "Recording MLD Join and Leave Events" on page 112.

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

disable

Syntax disable;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Disable MLD on the system.

Usage Guidelines See "Disabling MLD" on page 115.

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

exclude

Syntax exclude;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address*],
[edit protocols mld interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in JUNOS 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. By default, 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 Topics Enabling MLD Static Group Membership on page 105

group

Syntax

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

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

Release Information Statement introduced before JUNOS Release 7.4.

Description The MLD multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.

Options *multicast-group-address*—Address of the group.

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

The remaining statements are explained separately.

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

Related Topics Enabling MLD Static Group Membership on page 105

group-count

Syntax group-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address*],
[edit protocols mld interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in JUNOS Release 9.6.

Description Configure the number of static groups to be created.

Options *number*—Number of static groups.
Default: 1
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 Topics Enabling MLD Static Group Membership on page 105

group-increment

Syntax group-increment *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address* source],
[edit protocols mld interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in JUNOS Release 9.6.

Description Configure the number of times the address should be incremented for each static group created. The increment is specified in a format similar to an IPv6 address.

Options *increment*—Number of times the address should be incremented.
Default: ::1
Range: ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff:

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

Related Topics Enabling MLD Static Group Membership on page 105

group-policy

Syntax group-policy *policy-name*;

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

Release Information Statement introduced in JUNOS Release 9.1.

Description When this statement is enabled on a router running MLD version 1 (MLDv1), 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).

When this statement is enabled on a router running MLD version 2 (MLDv2), 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).

Usage Guidelines See "Filtering Unwanted MLD Reports at the MLD Interface Level" on page 103.

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

immediate-leave

Syntax immediate-leave;

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

Release Information Statement introduced in JUNOS Release 8.3.

Description When this statement is enabled on a router running MLDv1, after the router receives a multicast listener done message from a host associated with the interface, the router immediately removes the group membership from the interface and suppresses the sending of any group-specific queries for the multicast group.

When this statement is enabled on a router running MLDv2, after the router receives a report with the type BLOCK_OLD_SOURCES, the router suppresses the sending of group-and-source queries but relies on the JUNOS-supported host tracking mechanism to determine whether or not it removes a particular source group membership from the interface.

NOTE: Use this statement only on MLD interfaces to which one MLD host is connected. If more than one MLD host is connected to a LAN through the same interface, and one host sends a done message, the router removes all hosts on the interface from the multicast group. The router loses contact with the hosts that properly remain in the multicast group until they send join requests in response to the router's next general multicast listener query.

Usage Guidelines See "Specifying Immediate-Leave Host Removal for MLD" on page 102.

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

interface

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Enable MLD 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**. For details about specifying interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

The remaining statements are explained separately.

Usage Guidelines

See "Enabling MLD" on page 100.

Required Privilege Level

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

maximum-transmit-rate

Syntax maximum-transmit-rate *packets-per-second*;

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

Release Information Statement introduced in JUNOS Release 9.3.

Description Limit the transmission rate of MLD packets.

Options packets-per-second—Maximum number of MLD 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 Topics
- Limiting the Maximum MLD Message Rate on page 105

mld

Syntax

```
mld {
   accounting;
   interface interface-name {
      disable;
      (accounting | no-accounting);
      group-policy;
      immediate-leave;
      oif-map;
      passive;
      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>;
   }
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Enable MLD on the router. MLD must be enabled for the router to receive multicast packets.

Default

MLD is disabled on the router. MLD is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Options

The statements are explained separately.

Usage Guidelines

See “Enabling MLD” on page 100.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

no-accounting

See accounting.

oif-map

Syntax oif-map *map-name*;

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

Release Information Statement introduced in JUNOS Release 9.6.

Description Associate an outgoing interface (OIF) map to an MLD logical 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 Topics
- Configuring Multicast Outgoing Interface Mapping on page 220

passive (MLD)

Syntax passive;

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

Release Information Statement introduced in JUNOS Release 9.6.
allow-receive, send-general-query, and send-group-query options were added in JUNOS Release 10.0.

Description Specify that MLD run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as MLD reports, queries, and leaves.

NOTE: You can selectively activate up to two out of the three available options for the passive statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the passive statement.

Options allow-receive—Enables IGMP to receive control traffic on the interface.

send-general-query—Enables IGMP to send general queries on the interface.

send-group-query—Enables IGMP to send group-specific and group-source-specific queries on the interface.

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

Related Topics
- Configuring Multicast Outgoing Interface Mapping on page 220

query-interval

Syntax query-interval *seconds*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description How often the querier router sends general host-query messages.

Options *seconds*—Time interval.
Range: 1 through 1024
Default: 125 seconds

Usage Guidelines See "Modifying the MLD Host-Query Message Interval" on page 101.

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

Related Topics
- query-last-member-interval
- query-response-interval

query-last-member-interval

Syntax query-last-member-interval *seconds*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description How often the querier router 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 from 1 through 1024
Default: 1 second

Usage Guidelines See "Modifying the MLD Last-Member Query Interval" on page 102.

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

Related Topics
- query-interval
- query-response-interval

query-response-interval

Syntax query-response-interval *seconds*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description How long the querier router waits to receive a response to a host-query message from a host.

Options *seconds*—Time interval. This interval must be less than the interval between general host-query messages.
Range: 1 through 1024
Default: 10 seconds

Usage Guidelines See "Modifying the MLD Query Response Interval" on page 101.

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

Related Topics
- query-interval
- query-last-member-interval

robust-count

Syntax robust-count *number*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Tune for the expected packet loss on a subnet.

Options *number*—Time interval. This interval must be less than the interval between general host-query messages.
Range: 2 through 10
Default: 2 seconds

Usage Guidelines See "Modifying the MLD Robustness Variable" on page 104.

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

source

Syntax

```
source ip-address {
   source-count number;
   source-increment increment;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address*],
[edit protocols mld interface *interface-name* static group *multicast-group-address*]

Release Information

Statement introduced before JUNOS Release 7.4.

Description

IP version 6 (IPv6) unicast source address for the multicast group being statically configured on an interface.

Options

ip-address—One or more IPv6 unicast addresses.

Usage Guidelines

"Enabling MLD Static Group Membership" on page 105

Required Privilege Level

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

source-count

Syntax

source-count *number*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address* source],
[edit protocols mld interface *interface-name* static group *multicast-group-address* source]

Release Information

Statement introduced in JUNOS Release 9.6.

Description

Configure the number of multicast source addresses that should be accepted for each static group created.

Options

number—Number of source addresses.

Default: 1

Range: 1 through 1024

Required Privilege Level

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

Related Topics

Enabling MLD Static Group Membership on page 105

source-increment

Syntax source-increment *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address* source],
[edit protocols mld interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in JUNOS Release 9.6.

Description Configure the number of times the address should be incremented for each static group created. The increment is specified in a format similar to an IPv6 address.

Options *increment*—Number of times the source address should be incremented.
Default: ::1
Range: ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff:

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

Related Topics Enabling MLD Static Group Membership on page 105

ssm-map

Syntax ssm-map *ssm-map-name*;

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

Release Information Statement introduced in JUNOS Release 7.4.

Description Apply an SSM map to an MLD interface.

Options *ssm-map-name*—Name of SSM map.

Usage Guidelines See “Example: Configuring SSM Mapping” on page 208.

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

static

Syntax

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

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Test multicast forwarding on an interface.

The remaining statements are explained separately.

Usage Guidelines See "Enabling MLD Static Group Membership" on page 105.

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.

traceoptions

Syntax

```
traceoptions {
   file filename <files number> <size size> <world-readable | no-world-readable>;
   flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Configure MLD tracing options.

To specify more than one tracing operation, include multiple flag statements.

To trace the paths of multicast packets, use the mtrace command as described in the *JUNOS System Basics and Services Command Reference*.

Default

The default MLD trace options are those inherited from the traceoptions statement included at the [edit routing-options] hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place tracing output in the file mld-log.

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.
Range: 2 through 1000 files
Default: 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.

MLD Tracing Flags

- leave—Leave group messages.
- mtrace—Mtrace packets. Use the mtrace command to troubleshoot the software.
- packets—All MLD packets.

- query—MLD 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—Traces errors and significant events during normal packet processing

  Default: If you do not specify this option, only unusual or abnormal operations are traced.
- policy—Policy operations and actions
- route—Routing table changes
- state—State transitions
- task—Interface transactions and processing
- timer—Timer usage

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

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

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.
Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

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

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

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

If you specify a maximum file size, you must also include the files statement to specify the maximum number of trace files.
Syntax: *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB
Range: 10 KB through the maximum file size supported on your system
Default: 1 MB

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

Usage Guidelines See "Tracing MLD Protocol Traffic" on page 114.

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.

version

Syntax version *version*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the MLD version explicitly. MLD version 2 (MLDv2) is used only to support source-specific multicast (SSM).

Options *version*—MLD version to run on the interface.
Range: 1 or 2
Default: 1 (MLDv1)

Usage Guidelines See "Modifying the MLD Version" on page 101.

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.

Part 5
SAP and SDP

- SAP Overview on page 139
- SAP Configuration Guidelines Introduction on page 141
- Summary of SAP Configuration Statements on page 143

Chapter 11

SAP Overview

Session announcements are handled by two protocols: the Session Announcement Protocol (SAP) and the Session Description Protocol (SDP). These two protocols display multicast session names and correlate the names with multicast traffic. Only SAP has configuration parameters that users can change.

SDP is a session directory protocol that is used for multimedia sessions. It helps advertise multimedia conference sessions and communicates setup information to participants who want to join the session. SDP simply formats the session description; it does not incorporate a transport protocol. 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.

SAP is a session directory announcement protocol that SDP uses as its transport protocol.

For information about supported standards for SAP and SDP, see "JUNOS Software Multicast Standards" on page 17.

Chapter 12

SAP Configuration Guidelines Introduction

SAP Configuration Guidelines

The Session Announcement Protocol (SAP) and Session Description Protocol (SDP) associate multicast session names with multicast traffic addresses. Only SAP has configuration parameters that users can change. Enabling SAP allows the router to receive announcements about multimedia and other multicast sessions. To enable SAP and the receipt of session announcements, include the sap statement:

```
sap {
   disable;
   listen address <port port>;
}
```

You can include this statement at the following hierarchy levels:

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

For an overview of logical systems and a detailed example of logical system configuration, see the logical systems chapter of the *JUNOS Feature Guide*.

By default, SAP always listens to the address and port 224.2.127.254:9875 for session advertisements. To add other addresses or pairs of address and port, include one or more listen statements.

Sessions learned by SDP, SAP's higher-layer protocol, time out after 60 minutes.

Chapter 13

Summary of SAP Configuration Statements

The following sections explain each of the SAP multicast configuration statements. The statements are organized alphabetically.

disable

Syntax disable;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Explicitly disable SAP.

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

Related Topics

- SAP Configuration Guidelines on page 141

listen

Syntax listen *address* <port *port*>;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify an address and optionally a port on which SAP and SDP listen, in addition to the default SAP address and port on which they always listen, 224.2.127.254:9875. To specify multiple additional addresses or pairs of address and port, include multiple listen statements.

Options *address*—(Optional) Address on which SAP listens for session advertisements.
Default: 224.2.127.254

port *port*—(Optional) Port on which SAP listens for session advertisements.
Default: 9875

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

Related Topics
- SAP Configuration Guidelines on page 141

sap

Syntax

```
sap {
   disable;
   listen address <port port>;
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Enable the router to listen to session directory announcements for multimedia and other multicast sessions.

SAP and SDP always listen on the default SAP address and port, 224.2.127.254:9875. To have SIP listen on additional addresses or pairs of address and port, include a listen statement for each address or pair.

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 Topics

- listen
- SAP Configuration Guidelines on page 141

Part 6
PGM

- PGM Overview on page 149
- PGM Configuration Guidelines Overview on page 155
- Summary of PGM Configuration Statements on page 157

Chapter 14

PGM Overview

This chapter discusses the following topics:

- Pragmatic General Multicast Overview on page 149
- PGM Architecture and PGM Routers on page 150
- PGM-Enabled Source on page 151
- PGM-Enabled Receivers on page 151
- PGM-Enabled Routers on page 152

Pragmatic General Multicast Overview

Multicast applications often require real-time operation. These applications cannot take advantage of Transmission Control Protocol (TCP) reliability features such as sequencing, retransmission, and flow control through windowing between sender and receiver. The User Datagram Protocol (UDP), the major transport layer alternative to TCP, leaves much to be desired in its reliability for multicast traffic. Pragmatic General Multicast (PGM) is a 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 to request replacement information if the receiver application requires it. PGM is IP protocol number 113.

Although PGM is mainly concerned with the operation of multicast source and receiver, PGM-enabled routers (called PGM network elements) play a *router assistance* role in the initial delivery and potential replacement of multicast traffic. PGM routers are not mandatory in PGM, but they can provide the following benefits when placed anywhere between the source and receivers:

- Reduce the load on the multicast source by aggregating duplicate messages to the source. PGM routers are required to perform this function.
- Limit the flooding of *repair data* (replacement information) to only those downstream receivers that requested the repair data. PGM routers are required to perform this function.
- Act as *designated local repairers* (*DLRs*) by caching the repair data and resending it to receivers that request it later. DLR functions are a PGM option, and PGM routers are not required to perform this role.

PGM adds reliability to multicast traffic streams. It is not a complete multicast protocol like the Distance Vector Routing Multicast Protocol (DVMRP) or Protocol Independent

Multicast (PIM). Adding PGM to a router does not enable the router to perform multicast functions. Instead, a PGM router with multicast capabilities and a preconfigured multicast protocol such as PIM can offer more reliable multicast services to PGM sources and receivers. PGM is not an alternative to multicast routing protocols, but an enhancement of the multicast capabilities already present and configured on the router.

For information about supported standards for PGM, see "JUNOS Software Multicast Standards" on page 17.

PGM Architecture and PGM Routers

PGM is defined in RFC 3208 and forms a reliable transport layer for multicast applications. Almost any multicast application can use PGM. Applications most suitable for PGM include stock market ticker update information, news reports, weather warnings, and other information that must reach multiple listeners in its entirety and in a timely fashion.

The basic PGM architecture consists of a multicast content source, one or more receivers, and zero or more routers between source and receivers. All end devices must be PGM-enabled, although there can be non-PGM routers between source and receiver. If all routers are non-PGM routers, then no routers are capable of the PGM router assistance function, and all PGM functions take place directly between source and receiver.

PGM sources send sequenced content in *sessions* to receivers, using multicast protocols. Other, non-PGM protocols allow receivers to learn about a particular source, its sessions, and its location. PGM receivers listen to multicast *original data* (*ODATA*), detect missing content through the sequence numbers, and send unicast *negative acknowledgments* (*NAKs*) back to the source. NAKs are answered by multicast *NAK confirmations* (*NCFs*), which suppress any NAKs from receivers on the same subnet that have not yet sent a NAK upstream. The source sends multicast *repair data* (*RDATA*) to receivers containing the missing content. PGM routers assist in this process by making sure that the negative acknowledgments follow the same path as the outbound content upstream to the source, and by suppressing duplicate negative acknowledgments and repair information.

PGM sources must maintain a sliding window of retransmittable information. There is no concept of group membership in PGM, so receivers never need to communicate with the source unless they request repair data with a negative acknowledgment. However, this means that the PGM source determines the window size for each receiver, in contrast to almost all other protocols, and requires a certain processing power in each receiver. The absence of positive receiver-to-source acknowledgments also means that PGM scales well and cuts down on control message traffic that can easily overwhelm a multicast network.

PGM receivers can start receiving a PGM session from a PGM source at any time and request any missing previous information that the receiving application needs. If the session is long enough or the transmit window small enough so that the source does not maintain a long session history, the receiver cannot get all required information.

This section describes in more detail the behavior of the three PGM elements in a multicast network:

- PGM-Enabled Source on page 151
- PGM-Enabled Receivers on page 151
- PGM-Enabled Routers on page 152

PGM-Enabled Source

A PGM-enabled source of multicast content generates sequenced packets of ODATA that are multicast to receivers. Interleaved with the content packets are *source path messages* (*SPMs*), which tell PGM routers and receivers about their upstream next-hop PGM device—either another PGM router or the PGM source.

ODATA packets and SPMs are multicast from the source. A PGM router always appends its own IP address to the SPM before it is multicast on the downstream interfaces. The SPMs are sent by the source and upstream PGM routers with the router alert option set in the IP headers so that PGM routers do not have to examine every packet in the session for SPM packets.

The PGM source acknowledges a received NAK by multicasting an NCF downstream to the next PGM device on the path to the receiver. NCFs make sure that PGM routers and receivers do not bombard sources with NAKs. Downstream PGM routers suppress all subsequent NAKs that indicate the same missing information once one NCF is received from the upstream device.

The PGM source also responds to NAKs by multicasting RDATA packets with the same sequence number as the one indicated by the NAK. RDATA packets have the router alert option set in the IP header so that PGM routers can distinguish them from ODATA packets.

PGM sources organize their packets in sessions. PGM sources are not required to retain copies of information older than the current session, although they might. Long sessions are not necessarily kept on the source in their entirety.

PGM sources identify themselves through a global source ID (GSID). This globally unique source identifier is formed from the low-order 48 bits of the Message Digest 5 (MD5) signature of the Domain Name System (DNS) name of the source.

PGM-Enabled Receivers

The PGM architecture requires one or more PGM-enabled receivers of the multicast content generated by a PGM source. PGM receivers accept all types of downstream PGM messages: ODATA, SPMs, NCFs, and RDATA.

Receivers process the ODATA packets as they arrive from the source, constantly checking the 32-bit sequence number in the ODATA PGM header for gaps in the sequence. If the receiver detects missing information, it generates a NAK for that sequence number. The NAK is unicast upstream to the PGM next hop, which is a router or the source, as determined by the last address in the received SPM.

A receiver knows that its NAK was received by the PGM next hop when it gets an NCF in response to its NAK. If several receivers on a subnet are missing the same ODATA packet, receivers getting an NCF for the packet before sending a NAK suppress the NAK. If a receiver does not get an NCF in response to a NAK, the receiving application can send a NAK again or continue, with the certainty that information is missing.

After the NCF, PGM receivers get an RDATA packet with the same sequence number indicated in the NAK and a copy of the missing ODATA. NCFs and RDATA can originate from the source or a router acting as a DLR for a subnet. The receiver now has complete information or knows for certain what is missing.

PGM receivers can request almost anything from the PGM source. However, because the source determines the window size, there is no guarantee that older information is available.

PGM-Enabled Routers

Multicast-capable routers can implement the PGM router assistance functions, although not all multicast routers must be PGM-enabled routers. Mandatory PGM router assistance functions include aggregating duplicate NAKs sent to the source to reduce the load on the multicast source, generating NCFs in response to NAKs, and flooding RDATA packets to only those downstream receivers that requested it with a NAK. Optionally, a PGM router can be a DLR, caching PGM information and cutting down on network traffic by resending RDATA packets locally.

There can be zero or more PGM-enabled network elements (routers) between the source and receiver. If there are no PGM routers between the source and receiver, then all PGM messages flow directly between the source and receiver, and no router assistance functions are possible. Both PGM and non-PGM routers can be freely mixed on a network because PGM is a transport layer protocol and is not involved with router multicast functions.

PGM routers also receive SPMs from the source or an upstream PGM router and forward them downstream, inserting the router's own downstream IP interface address into the SPM so that receivers always know their upstream PGM next hop.

When a PGM router receives unicast NAKs from a downstream PGM router or receiver, the router unicasts one NAK for each missing sequence number to the next-hop PGM device upstream toward the source. The address of the PGM next-hop device is determined by received SPMs.

The PGM router multicasts NCFs in response to received NAKs on the downstream interfaces that received the NAKs. NCFs are not multicast on interfaces that have not received NAKs.

PGM routers must multicast all ODATA and RDATA packets they receive from upstream PGM devices. Normal multicast protocols are used to determine downstream interfaces.

If the PGM router is a DLR, it responds to received NAKs with an NCF and with its own RDATA packet. NAKs are not forwarded upstream from a DLR.

Figure 10 on page 153 shows the overall PGM architecture and the role of PGM-enabled routers.

Figure 10: PGM Architecture and General Operation

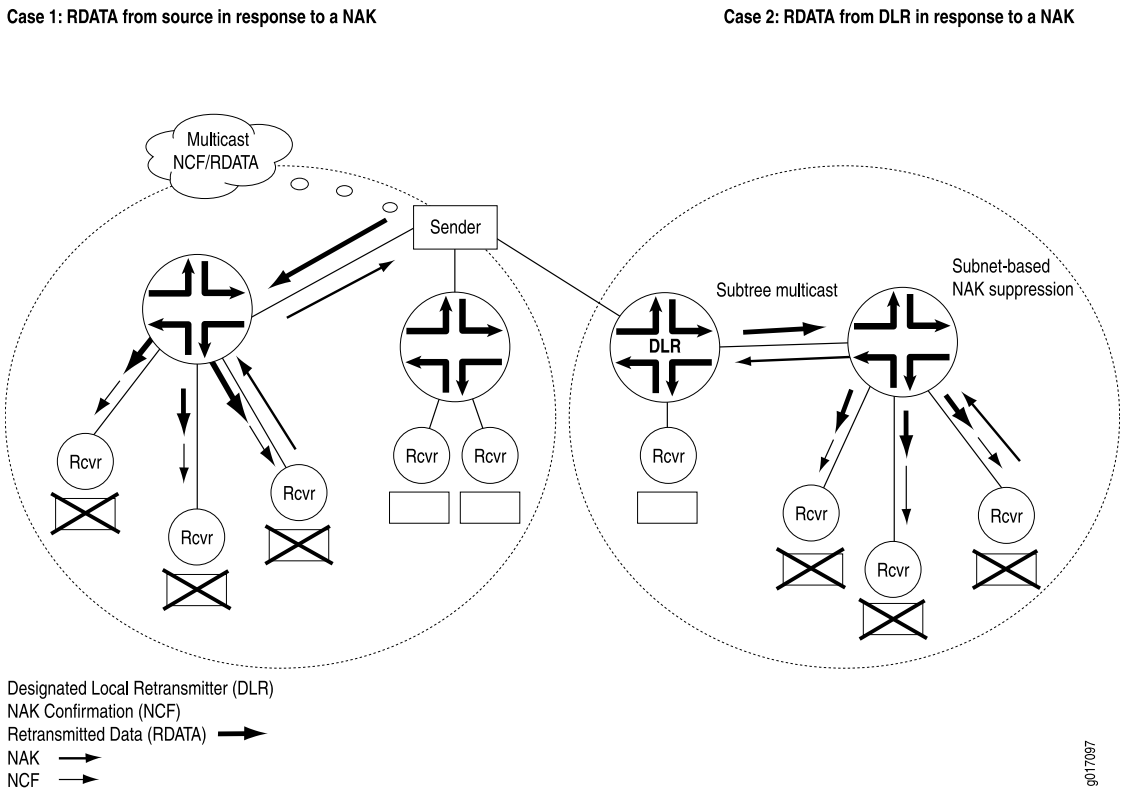

The figure shows only NAKs, NCFs, and RDATA flows. RDATA can come from either the source (left) or a DLR router (right). In both cases, unicast NAKs from a receiver are forwarded upstream by the routers, and multicast NCFs are generated downstream. Subnet NAK suppression is shown, as well as RDATA from the source or DLR sent only to the portions of the network requesting it.

Chapter 15

PGM Configuration Guidelines Overview

PGM Configuration Guidelines

Pragmatic General Multicast (PGM) allows the router to participate in defined PGM router assistance functions between PGM-enabled sources and receivers. Although PGM is a transport layer protocol and is not directly concerned with IP packet routing, PGM must be explicitly configured on the router.

To enable PGM globally on the router, include the pgm statement:

```
pgm;
```

You can include this statement at the following hierarchy levels:

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

To trace the operation of PGM, include the traceoptions statement:

```
traceoptions {
    flag flag <flag-modifier>;
}
```

You can include this statement at the following hierarchy levels:

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

You can specify the following PGM-specific options in the flag statement:

- all—Trace all PGM packets.
- init—Trace all PGM initialization events.
- packets—Trace all PGM packet processing.
- parser—Trace all PGM parser processing.
- route-socket—Trace all PGM route-socket events.
- show—Trace all PGM show command servicing.
- state—Trace all PGM state transitions.

For an overview of logical systems and a detailed example of logical system configuration, see the logical systems chapter of the *JUNOS Feature Guide*.

By default, PGM is enabled on every interface of the router, but global, explicit configuration is required. No options are available for PGM operation.

Chapter 16

Summary of PGM Configuration Statements

The following sections explain each PGM configuration statement. The statements are organized alphabetically.

pgm

Syntax

```
pgm {
   traceoptions {
      flag flag <flag-modifier>;
   }
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Configure PGM globally and set tracing options.

To specify more than one tracing operation, include multiple flag statements.

The remaining statement is explained separately.

Default

The default PGM trace options are inherited from the routing protocol traceoptions statement included 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 Topics

- PGM Configuration Guidelines on page 155

traceoptions

Syntax

```
traceoptions {
   flag flag <flag-modifier>;
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Configure PGM tracing options.

To specify more than one tracing operation, include multiple flag statements.

Default

The default PGM trace options are those inherited from the routing protocol traceoptions statement included at the [edit routing-options] hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.

PGM Tracing Flags

- all—Trace all PGM packets.
- init—Trace all PGM initialization events.
- packets—Trace all PGM packet processing.
- parser—Trace all PGM parser processing.
- route-socket—Trace all PGM route-socket events.
- show—Trace all PGM show command servicing.
- state—Trace all PGM state transitions.

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 the following modifiers:

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

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 Topics
- PGM Configuration Guidelines on page 155

Part 7
Multicast Routing Instances

- Multicast over Layer 3 Draft-Rosen VPNs Overview on page 163
- Draft-Rosen Multicast VPN Configuration Guidelines on page 169
- Summary of Draft-Rosen MVPN Configuration Statements on page 177

Chapter 17

Multicast over Layer 3 Draft-Rosen VPNs Overview

This chapter provides the following information about multicast over draft-rosen virtual private networks (VPNs) with service provider tunnels operating in source-specific mode (SSM):

- Draft-Rosen Multicast VPNs Overview on page 163
- Data MDTs Overview on page 164
- Data MDT Characteristics on page 165
- Draft-Rosen Multicast VPN Autodiscovery on page 166
- Draft-Rosen Multicast VPN Control Plane on page 166
- Interoperating with Other Vendors on page 167

Draft-Rosen Multicast VPNs Overview

The JUNOS Software provides three ways to configure IP version 4 (IPv4) multicast over Layer 3 VPNs:

- Draft-rosen multicast VPNs—Described in RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)* and based on the IETF Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP IP VPNs*. Draft-rosen multicast VPNs can be configured with service provider tunnels operating in any-source multicast (ASM) mode or source-specific multicast (SSM) mode. For more information about these types of draft-rosen multicast VPNs, see the *JUNOS Feature Guide*.
- Draft-rosen multicast VPNs—Described in RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)* and based on Section Two of the IETF Internet draft draft-rosen-vpn-mcast-06.txt, *Multicast in MPLS/BGP VPNs* (expires April 2004).
- Multiprotocol BGP-based multicast VPNs (also referred to as next-generation Layer 3 VPN multicast)—Defined by a set of sender sites and a set of receiver sites, they use BGP as the signaling protocol. For more information about multiprotocol BGP-based multicast VPNs, see the *JUNOS VPNs Configuration Guide*.

Draft-rosen multicast VPNs (MVPNs) with service provider tunnels operating in SSM mode do not require that the provider (P) routers maintain any VPN-specific PIM information.

Draft-rosen MVPNs with service provider tunnels operating in SSM mode, make use of two types of multicast distribution trees (MDTs); a *default MDT* and one or more *data MDTs*.

Draft-rosen MVPNs with service provider tunnels operating in SSM mode start by sending all multicast traffic over a default MDT, as described in section 7 of Internet draft draft-rosen-vpn-mcast-07.txt. This default mapping results in the delivery of packets to each provider edge (PE) router attached to the provider router even if the PE router has no receivers for the multicast group in that VPN. Each PE router processes the encapsulated VPN traffic even if the multicast packets are then discarded.

A data MDT, as described in "Data MDTs Overview" on page 164, solves the problem of routers flooding unnecessary multicast information to PE routers that have no interested receivers for a particular VPN multicast group.

The default MDT uses multicast tunnel (mt-) logical interfaces. Data MDTs also use mt logical interfaces. If you administratively disable the physical interface that the mt logical interfaces are configured on, the mt logical interfaces are moved to a different physical interface that is up. In this case the traffic is sent over the default MDT until new data MDTs are created as described in "Data MDT Characteristics" on page 165.

Data MDTs Overview

For draft-rosen MVPNs with service provider tunnels operating in SSM mode, the PE routers initially discover one another in a VPN routing and forwarding (VRF) instance across the default MDT. Any PE router configured with a default MDT group address can be the multicast source of one default MDT.

The PE routers can be configured so that when the multicast source within a site exceeds a traffic rate threshold, the PE router to which the source site is attached creates a new data MDT and advertises a new MDT group address.

The advertisement is sent in a User Datagram Protocol (UDP) type length value (TLV) packet called an *MDT join TLV*. The MDT join TLV identifies the source and group pair (S,G) in the VRF instance and the new data MDT group address used in the provider space. The PE router to which the source site is attached periodically announces the MDT join TLV over the default MDT for that VRF instance as long as the source is active. All PE routers receive the MDT join TLV because it is sent over the default MDT.

The traffic from that source is sent over the new data MDT. PE routers that have interested receivers can join the multicast group and receive the stream over the new data MDT. PE routers that do not have interested receivers do not join the group. After the PE routers join the new data MDT, the multicast traffic for that group is no longer sent over the default MDT.

Only the PE routers with receivers in the VRF instance for that multicast group join the new group. The PE routers must join the new group to receive the multicast traffic now sent over the new MDT by the source PE. PE routers without interested receivers listed in the VRF instance ignore the MDT join TLVs.

When remote PE routers join the new data MDT group, they send a PIM join message for the new group. The PIM join message for the new group is sent directly to the source PE router from the remote PE routers by means of PIM source-specific multicast. SSM mode using (S,G) is possible with a data MDT because the source address is known from the signaling used with a data MDT.

The source PE router starts encapsulating the multicast traffic for the VRF instance using the new data MDT group after 3 seconds, allowing time for the remote PE routers to join the new group. The source PE router then halts the flow of multicast packets over the default MDT, and the packet flow for the VRF instance source shifts to the newly created data MDT.

The PE router monitors the traffic rate during its periodic statistics-collection cycles. If the traffic rate drops below the threshold or the source stops sending, the PE router to which the source site is attached stops announcing the MDT join TLVs and switches back to sending on the default MDT for that VRF instance.

Data MDT Characteristics

The maximum number of data MDTs for all VPNs on a PE router is 8000, and the maximum number of data MDTs for a VRF instance is 1024. The configuration of a VRF instance can limit the number of MDTs possible. No new MDTs can be created after the 1024 MDT limit is reached in the VRF instance, and all traffic for other sources that exceed the configured limit is sent on the default MDT.

Creation of data MDTs depends on the monitoring of the multicast source data rate. This rate is checked once per minute, so the creation of data MDTs can be delayed up to 1 minute after a source exceeds a configured limit. In the same way, if the source data rate falls below the configured value, data MDT deletion can be delayed for up to 1 minute until the next statistics-monitoring collection cycle.

Changes to the configured data MDT limit value do not affect existing tunnels that exceed the new limit. Data MDTs that are already active remain in place until the threshold conditions are no longer met.

To remove active data MDTs no longer included in a newly configured group address range, you must restart the PIM routing instance. This restart clears all remnants of the former group addresses but disrupts routing and therefore requires a maintenance window for the change.

Multicast tunnel (mt-) interfaces created because of exceeded thresholds are not re-created if the routing process crashes. Therefore, graceful restart does not automatically reinstate the data MDT state. However, as soon as the periodic statistics collection reveals that the threshold condition is still exceeded, the tunnels are quickly re-created.

Draft-Rosen Multicast VPN Autodiscovery

A draft-rosen MVPN with service provider tunnels operating in SSM mode uses BGP signaling for autodiscovery of the PE routers.

Each PE sends an MDT subsequent address family identifier (MDT-SAFI) BGP network layer reachability information (NLRI) advertisement. The advertisement contains the following:

- Route distinguisher
- Unicast address of the PE router to which the source site is attached (usually the loopback)
- Multicast group address
- Route target extended community attribute

Each remote PE router imports the MDT-SAFI advertisements from each of the other PE routers if the route target matches. Each PE router then joins the (S,G) tree rooted at each of the other PE routers.

After a PE router discovers the other PE routers, the source and group are bound to the VRF through the multicast tunnel de-encapsulation interface.

A draft-rosen MVPN with service provider tunnels operating in any-source multicast sparse-mode uses a shared tree and rendezvous point (RP) for autodiscovery of the PE routers. The PE that is the source of the multicast group encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router. The RP then builds a shortest-path tree (SPT) toward the source PE. The remote PE that acts as a receiver for the MDT multicast group sends (*, G) join messages toward the RP and joins the distribution tree for that group.

Draft-Rosen Multicast VPN Control Plane

The control plane of a draft-rosen MVPN with service provider tunnels operating in SSM mode must be configured to support autodiscovery.

After the PE routers are discovered, PIM is notified of the multicast source and group addresses. PIM binds the (S,G) state to the multicast tunnel (mt) interface and sends a join message for that group.

Autodiscovery for a draft-rosen MVPN with service provider tunnels operating in SSM mode uses some of the facilities of the BGP-based MVPN control plane software module. Therefore, the BGP-based MVPN control plane must be enabled. The BGP-based MVPN control plane can be enabled for autodiscovery only.

Interoperating with Other Vendors

If you are configuring your Juniper Networks router to interoperate with other vendors' routers, verify the following:

- All provider tunnels use the same group address.
- All routers are configured with the same IP address for the lo0.*n* interface in the routing instance as the IP address for the lo0.0 interface in the master instance.
- The other vendor's router supports route target extended community attributes.
- The other vendor's router is configured to export route target extended community attributes.

Chapter 18

Draft-Rosen Multicast VPN Configuration Guidelines

This chapter provides the following information about the configuration of draft-rosen multicast VPNs:

- Configuring Draft-Rosen Multicast VPNs on page 169
- Configuring Data MDTs and Provider Tunnels on page 169
- Enabling Dynamic Reuse of Data MDT Group Addresses on page 173
- Configuring the Provider Tunnel SSM Group Address on page 173
- Configuring Draft-Rosen Multicast VPN Autodiscovery on page 173
- Controlling the Single Forwarder Election on page 175
- Configuring the Draft-Rosen Multicast VPN Control Plane on page 175
- Configuring Multicast for Virtual Routers on page 176

Configuring Draft-Rosen Multicast VPNs

Draft-rosen multicast virtual private networks (MVPNs) can be configured to support service provider tunnels operating in any-source multicast (ASM) mode or source-specific multicast (SSM) mode.

There are three functional areas to configure for draft-rosen MVPNs with service provider tunnels operating in SSM mode:

- Provider tunnel
- Autodiscovery mechanism
- Control plane protocol

Users need to be familiar with Layer 3 VPN operation on Juniper Networks routers, which is discussed in the JUNOS VPNs Configuration Guide.

Configuring Data MDTs and Provider Tunnels

To configure data multicast distribution trees (MDTs), include the mdt statement:

```
mdt {
    data-mdt-reuse
```

```
    group-range multicast-prefix;
    threshold {
        group group-address {
            source source-address {
                rate threshold-rate;
            }
        }
    }
    tunnel-limit limit;
}
```

You can include the statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols pim]
- [edit routing-instances *routing-instance-name* provider-tunnel]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel]

Data MDTs configured at the [edit routing-instances *routing-instance-name* provider-tunnel] hierarchy level support service provider tunnels operating in source-specific multicast mode. Data MDTs configured at the [edit routing-instances *routing-instance-name* protocols pim] hierarchy level support service provider tunnels operating in any-source multicast mode.

For an overview of routing instances and a detailed example of routing instance configuration, see the Multicast over Layer 3 VPNs chapter of the *JUNOS Feature Guide*.

By default, automatic creation of data MDTs is disabled.

This section describes the following tasks for configuring data MDTs:

- Configuring the Data MDT Group Range on page 170
- Configuring the Data MDT Threshold Parameters on page 171
- Limiting the Maximum Number of Data MDTs on page 172

Data MDTs require a correctly configured Layer 3 VPN for multicast. For more information about configuring Layer 3 VPNs for multicast, see “Configuring Any-Source Multicast for Draft-Rosen VPNs” on page 415.

Configuring the Data MDT Group Range

The provider edge (PE) router implementing data MDTs for a local multicast source must be configured with a range of multicast group addresses. Group addresses that fall within the configured range are used in the join messages for the data MDTs created in this VRF instance. The group address range cannot overlap the default MDT group address configured for any VPN on the router. If you configure overlapping group addresses, the configuration commit operation fails.

To configure the data MDT group range, include the **group-range** statement:

```
group-range multicast-prefix;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols pim mdt]
- [edit routing-instances *routing-instance-name* provider-tunnel mdt]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim mdt]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel mdt]

Any multicast address range can be used as the multicast prefix, for example, 227.0.0.0/8.

Configuring the Data MDT Threshold Parameters

The PE router implementing data MDTs for a local multicast source must establish a data MDT-creation threshold for a multicast group and source. A multicast group can have more than one source of traffic.

To configure the data MDT-creation threshold, include the **threshold** statement:

```
threshold {
    group group-address {
        source source-address {
            rate threshold-rate;
        }
    }
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols pim mdt]
- [edit routing-instances *routing-instance-name* provider-tunnel mdt]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim mdt]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel mdt]

The group address is the multicast group address to which the threshold applies. This could be a well-known address for a certain type of multicast traffic.

The source address is the unicast address of the source of the multicast traffic. It can be a source locally attached to or reached through the PE router. A group can have more than one source.

The group and source addresses can be explicit (all 32 bits of the address specified) or a prefix (network address and prefix length specified). Explicit and prefix address

forms can be combined if they do not overlap. Overlapping configurations, in which prefix and more explicit address forms are used for the same source or group address, are not supported.

The rate is the threshold applied to the multicast source to automatically create a data MDT. The range is from 10 kilobits per second (Kbps), the default, to 1 gigabit per second (Gbps, equivalent to 1,000,000 Kbps).

Limiting the Maximum Number of Data MDTs

The PE router implementing a data MDT for a local multicast source must establish a limit for the number of data MDTs created in this VRF instance. If the limit is 0 (the default), then no data MDTs are created for this VRF instance.

To configure the data MDT limit, include the tunnel-limit statement:

```
tunnel-limit limit;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols pim mdt]
- [edit routing-instances *routing-instance-name* provider-tunnel mdt]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim mdt]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel mdt]

The valid range is from 0 through 1024 for a VRF instance. There is a limit of 8000 tunnels for all data MDTs in all VRF instances on a PE router.

M Series and T Series routers that are connected to a source site require a Tunnel Services PIC to support multicast tunnel (mt) interfaces. MX Series routers do not require Tunnel Services PICs.

A Tunnel Services PIC supports a maximum of 512 multicast tunnel interfaces. Both default and data MDTs contribute to this total. The default MDT uses two multicast tunnel interfaces (one for encapsulation and one for de-encapsulation). To enable an M Series or T Series router to support more than 512 multicast tunnel interfaces, another Tunnel Services PIC is required.

For more information about Tunnel Services PICs and multicast tunnels, see “Tunnel Services PICs and Multicast” on page 362.

Enabling Dynamic Reuse of Data MDT Group Addresses

This section describes how to enable dynamic reuse of data multicast distribution tree (MDT) group addresses.

A limited number of multicast group addresses are available for use in data MDT tunnels. By default when the available multicast group addresses are all used, no new data MDTs can be created.

You can enable dynamic reuse of data MDT group addresses. Dynamic reuse of data MDT group addresses allows multiple multicast streams to share a single MDT and multicast provider group address. For example, three streams can use the same provider group address and MDT tunnel.

To enable dynamic reuse of data MDT group addresses, include the data-mdt-reuse statement:

```
user@host# set protocols pim mdt data-mdt-reuse
```

The streams are assigned to a particular MDT in a round-robin fashion. Since a provider tunnel might be used by multiple customer streams, this can result in egress routers receiving customer traffic that is not destined for their attached customer sites.

Related Topics

- data-mdt-reuse

Configuring the Provider Tunnel SSM Group Address

To configure the PIM-SSM provider tunnel default MDT group address, include the group-address statement:

```
group-address address;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel pim-ssm]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel pim-ssm]

NOTE: When you configure the default MDT group address by including the group-address statement, do not specify a default MDT group address by including the vpn-group-address statement.

Configuring Draft-Rosen Multicast VPN Autodiscovery

To configure draft-rosen VPN autodiscovery for provider tunnels operating in SSM mode, include the mvpn statement:

```
mvpn {
   autodiscovery {
      inet-mdt;
   }
}
```

You can include this statement at the following hierarchy levels:

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

To enable autodiscovery, include the inet-mdt statement.

NOTE: When you configure draft-rosen multicast VPNs with provider tunnels operating in source-specific mode and using the vrf-export statement to specify the export policy, the policy must have a term that accepts routes from the vrf-name.mdt.0 routing table. This term ensures proper PE autodiscovery using the inet-mdt address family.

When you configure draft-rosen multicast VPNs with provider tunnels operating in source-specific mode and using the vrf-target statement, the VRF export policy is automatically generated and automatically accepts routes from the vrf-name.mdt.0 routing table.

To configure the BGP-based MVPN control plane to provide signaling only for autodiscovery and not for PIM operations, include the autodiscovery-only statement:

```
autodiscovery-only {
   intra-as {
      inclusive;
   }
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn]

Related Topics

- Draft-Rosen Multicast VPNs Overview on page 163
- Data MDTs Overview on page 164
- Data MDT Characteristics on page 165
- Draft-Rosen Multicast VPN Autodiscovery on page 166
- Draft-Rosen Multicast VPN Control Plane on page 166
- Interoperating with Other Vendors on page 167

Controlling the Single Forwarder Election

When a PE receives a customer join or prune message from a CE, the message identifies a particular multicast flow as belonging either to a source-specific tree (S,G) or to a shared tree (*,G). If the route to the multicast source or RP is across the VPN backbone, then the PE needs to identify the upstream multicast hop (UMH) for the (S,G) or (*,G) flow. Normally the UMH is determined by the unicast route to the multicast source or RP.

However, in some cases, the CEs may be distributing to the PEs a special set of routes that are to be used exclusively for the purpose of upstream multicast hop selection using the route-import community. More than one route might be eligible, and the PE needs to elect a single forwarder from the eligible UMHs.

The default metric for the single forwarder election is the IP address from the global administrator field in the route-import community. You can configure a router to use the unicast route preference to determine the single forwarder election.

To configure the router to use the unicast route preference, include the unicast-umh-election statement.

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Related Topics

- unicast-umh-election

Configuring the Draft-Rosen Multicast VPN Control Plane

To enable MDT-SAFI NLRI control plane messages, include the signaling statement under the family statement and specify the inet-mdt option:

```
family inet-mdt {
    signaling;
}
```

You can include this statement at the following hierarchy levels:

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

You must enable the BGP-based MVPN control plane for draft-rosen MVPN autodiscovery. To enable the BGP-based MVPN control plane software module for autodiscovery only, include the autodiscovery-only statement:

```
autodiscovery-only {
    intra-as {
        inclusive;
    }
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn]

Configuring Multicast for Virtual Routers

You can configure PIM for the virtual-router routing instance type as well as for the vrf instance type. The virtual-router instance type is similar to the vrf instance type used with Layer 3 VPNs, but is used for non-VPN-related applications.

The virtual-router instance type has no VRF import, VRF export, VRF target, or route distinguisher requirements. The virtual-router instance type is used for non-Layer 3 VPN situations, for example, to allow the use of IP Security (IPsec) tunnels within VPNs.

When PIM is configured under the virtual-router instance type, the VPN configuration is not based on RFC 2547, *BGP/MPLS VPNs*, so PIM operation does not comply with the Internet draft draft-rosen-vpn-mcast-07.txt *Multicast in MPLS/BGP VPNs*. For more information about multicast draft support, see "IP Multicast Specifications" on page 18. In the virtual-router instance type, PIM operates in a routing instance by itself, forming adjacencies with PIM neighbors over the routing instance interfaces as the other routing protocols do with neighbors in the routing instance.

To configure PIM for a virtual-router instance type, include both the pim statement and the instance-type statement, specifying the virtual-router instance type:

```
instance-type virtual-router;
protocols {
    pim {
        ...pim-configuration...
    }
}
```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Do not include the vpn-group-address statement for the virtual-router instance type.

Chapter 19

Summary of Draft-Rosen MVPN Configuration Statements

The following sections explain draft-rosen MVPN configuration statements. The statements are organized alphabetically.

autodiscovery

Syntax

```
autodiscovery {
    inet-mdt;
}
```

Hierarchy Level

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

Release Information

Statement introduced in JUNOS Release 9.4.

Description

Enable the PE routers in the VPN to discover one another automatically.

Options

The statements are explained separately.

Usage Guidelines

See "Configuring Draft-Rosen Multicast VPN Autodiscovery" on page 173.

Required Privilege Level

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

autodiscovery-only

Syntax

```
autodiscovery-only {
    intra-as {
        inclusive;
    }
}
```

Hierarchy Level

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

Release Information

Statement introduced in JUNOS Release 9.4.

Description

Enable the BGP-based MVPN control plane for autodiscovery only.

Usage Guidelines

See "Configuring Draft-Rosen Multicast VPN Autodiscovery" on page 173.

Required Privilege Level

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

data-mdt-reuse

Syntax data-mdt-reuse;

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

Release Information Statement introduced in JUNOS Release 10.0.

Description Enable dynamic reuse of data MDT group addresses.

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

Related Topics
- Enabling Dynamic Reuse of Data MDT Group Addresses on page 173

group

Syntax

```
group group-address {
   source source-address {
      rate threshold-rate;
   }
}
```

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the explicit or prefix multicast group address to which the threshold limits apply. This is typically a well-known address for a certain type of multicast traffic.

Options *group-address*—Explicit group address to limit.

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Data MDT Threshold Parameters” on page 171.

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

group-address

Syntax group-address *address*;

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

Release Information Statement introduced in JUNOS Release 9.4.

Description Configure the PIM-SSM provider tunnel default MDT using the specified group address.

Usage Guidelines See "Configuring the Provider Tunnel SSM Group Address" on page 173.

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

group-range

Syntax group-range *multicast-prefix*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Establish the group range to use for data MDTs created in this VRF instance. This address range cannot overlap the default MDT addresses of any other VPNs on the router. If you configure overlapping group ranges, the configuration commit fails.

Options *multicast-prefix*—Multicast address range to identify data MDTs.
Range: Any valid, nonreserved multicast address range
Default: None

Usage Guidelines See "Configuring the Data MDT Group Range" on page 170.

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

inclusive

Syntax inclusive;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn autodiscovery-onlyintra-as],
[edit routing-instances *routing-instance-name* protocols mvpn autodiscovery-only intra-as]

Release Information Statement introduced in JUNOS Release 9.4.

Description Enable the MVPN control plane for autodiscovery only, using intra-AS autodiscovery routes over an inclusive provider multicast service interface (PMSI).

Usage Guidelines See "Configuring Draft-Rosen Multicast VPN Autodiscovery" on page 173.

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

inet-mdt

Syntax inet-mdt;

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

Release Information Statement introduced in JUNOS Release 9.4.

Description Configure the PE router in a VPN to use a multicast distribution tree (MDT) subsequent address family identifier (SAFI) NLRI for autodiscovery of other PE routers.

Usage Guidelines See "Configuring Draft-Rosen Multicast VPN Autodiscovery" on page 173.

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

intra-as

Syntax

```
intra-as {
    inclusive;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn autodiscovery-only],
[edit routing-instances *routing-instance-name* protocols mvpn autodiscovery-only]

Release Information

Statement introduced in JUNOS Release 9.4.

Description

Enable the MVPN control plane for autodiscovery only, using intra-AS autodiscovery routes.

Usage Guidelines

See "Configuring Draft-Rosen Multicast VPN Autodiscovery" on page 173.

Required Privilege Level

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

mdt

Syntax

```
mdt {
   data-mdt-reuse
   group-range multicast-prefix;
   threshold {
      group group-address {
         source source-address {
            rate threshold-rate;
         }
      }
      tunnel-limit limit;
   }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel],
[edit routing-instances *routing-instance-name* protocols pim],
[edit routing-instances *routing-instance-name* provider-tunnel]

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Establish the group address range for data MDTs, the threshold for the creation of data MDTs, and tunnel limits for a multicast group and source. A multicast group can have more than one source of traffic.

The remaining statements are explained separately.

Usage Guidelines

See “Configuring Data MDTs and Provider Tunnels” on page 169.

Required Privilege Level

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

mvpn (Control Plane Autodiscovery)

Syntax

```
mvpn {
   autodiscovery-only {
      intra-as {
         inclusive;
      }
   }
   receiver-site;
   route-target {
      export-target {
         target target-community;
         unicast;
      }
      import-target {
         target {
            target <target:number:number> <receiver | sender>;
            unicast <receiver | sender>;
         }
         unicast {
            receiver;
            sender;
         }
      }
   }
   sender-site;
   traceoptions {
      file filename <files number> <size maximum-file-size> <world-readable |
         no-world-readable>;
      flag flag <flag-modifier> <disable>;
   }
   unicast-umh-election;
}
```

Hierarchy Level

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

Release Information

Statement introduced in JUNOS Release 9.4.

Description

Enable the MVPN control plane for autodiscovery only.

Usage Guidelines

See “Configuring Draft-Rosen Multicast VPN Autodiscovery” on page 173.

Required Privilege Level

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

mvpn (Routing Instance)

Syntax

```
mvpn {
   autodiscovery {
      inet-mdt;
   }
}
```

Hierarchy Level

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

Release Information

Statement introduced in JUNOS Release 9.4.

Description

Configure the control plane to be used for PE routers in the VPN to discover one another automatically.

Options

The other statements are explained separately.

Usage Guidelines

See "Configuring Draft-Rosen Multicast VPN Autodiscovery" on page 173.

Required Privilege Level

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

pim-ssm

Syntax

```
pim-ssm {
   group-address address;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel],
[edit routing-instances *routing-instance-name* provider-tunnel]

Release Information

Statement introduced in JUNOS Release 9.4.

Description

Configure the PIM-SSM provider tunnel.

The remaining statement is explained separately.

Usage Guidelines

See "Configuring the Provider Tunnel SSM Group Address" on page 173.

Required Privilege Level

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

rate

Syntax rate *threshold-rate*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim mdt threshold group *group-address* source *source-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel mdt threshold group *group-address* source *source-address*],
[edit routing-instances *routing-instance-name* protocols pim mdt threshold group *group-address* source *source-address*],
[edit routing-instances *routing-instance-name* provider-tunnel mdt threshold group *group-address* source *source-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Apply a rate threshold to a multicast source to automatically create a data MDT.

Options *threshold-rate*—Rate in kilobits per second (Kbps) to apply to source.
Range: 10 Kbps through 1 Gbps (1,000,000 Kbps)
Default: 10 Kbps

Usage Guidelines See "Configuring the Data MDT Threshold Parameters" on page 171.

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

signaling

Syntax signaling;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp family inet-mdt],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp family inet-mvpn],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* family inet-mdt],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* family inet-mvpn],
[edit routing-instances *routing-instance-name* protocols bgp family inet-mdt],
[edit routing-instances *routing-instance-name* protocols bgp family inet-mvpn],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* family inet-mdt],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* family inet-mvpn]

Release Information Statement introduced in JUNOS Release 9.4.

Description Enable signaling in BGP. For multicast distribution tree (MDT) subaddress family identifier (SAFI) NLRI signaling, configure signaling under the inet-mdt family. For multiprotocol BGP (MBGP) intra-AS NLRI signaling, configure signaling under the inet-mvpn family.

Usage Guidelines See “Configuring the Draft-Rosen Multicast VPN Control Plane” on page 175.

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

source

Syntax

```
source source-address {
    rate threshold-rate;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim mdt threshold group *group-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel mdt threshold group *group-address*],
[edit routing-instances *routing-instance-name* protocols pim mdt threshold group *group-address*],
[edit routing-instances *routing-instance-name* provider-tunnel mdt threshold group *group-address*]

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Establish a threshold to trigger the automatic creation of a data MDT for the specified unicast address or prefix of the source of multicast information.

Options

source-address—Explicit unicast address of the multicast source.

The remaining statement is explained separately.

Usage Guidelines

See “Configuring the Data MDT Threshold Parameters” on page 171.

Required Privilege Level

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

threshold

Syntax

```
threshold {
   group group-address {
      source source-address {
         rate threshold-rate;
      }
   }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim mdt],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel mdt],
[edit routing-instances *routing-instance-name* protocols pim mdt],
[edit routing-instances *routing-instance-name* provider-tunnel mdt]

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Establish a threshold to trigger the automatic creation of a data MDT.

The remaining statements are explained separately.

Usage Guidelines

See "Configuring the Data MDT Threshold Parameters" on page 171.

Required Privilege Level

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

tunnel-limit

Syntax tunnel-limit *limit*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Limit the number of data MDTs created in this VRF instance. If the limit is 0, then no data MDTs are created for this VRF instance.

Options *limit*—Maximum number of data MDTs for this VRF instance.
Range: 0 through 1024
Default: 0 (No data MDTs are created for this VRF instance.)

Usage Guidelines See "Limiting the Maximum Number of Data MDTs" on page 172.

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

unicast-umh-election

Syntax unicast-umh-election;

Hierarchy Level [edit routing-instances *routing-instance-name* protocols mvpn]

Release Information Statement introduced in JUNOS Release 9.6.

Description Configure a router to use the unicast route preference to determine the single forwarder election.

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

Related Topics
- Controlling the Single Forwarder Election on page 175
- mvpn (Control Plane Autodiscovery)

Part 8
Multicast Routing Options

- Multicast Administrative Scoping Configuration Guidelines on page 193
- Multicast Reverse Path Forwarding Configuration Guidelines on page 199
- Source-Specific Multicast Configuration Guidelines on page 203
- Flow Map Configuration Guidelines on page 211
- Bandwidth Management Configuration Guidelines on page 215
- Multicast Forwarding Cache Configuration Guidelines on page 229
- Ingress PE Redundancy Configuration Guidelines on page 235
- Configuring PIM-to-IGMP and PIM-to-MLD Message Translation on page 239
- Summary of Multicast Routing Options Configuration Statements on page 245

Chapter 20

Multicast Administrative Scoping Configuration Guidelines

This chapter discusses the following topics:

- Multicast Scoping Overview on page 193
- Creating a Named Scope for Multicast Scoping on page 194
- Using a Scope Policy for Multicast Scoping on page 195
- Example: Creating a Named Scope for Multicast Scoping on page 196
- Example: Using a Scope Policy for Multicast Scoping on page 197

Multicast Scoping Overview

You use multicast scoping to limit multicast traffic by configuring it to an administratively defined topological region. Multicast scoping controls the propagation of multicast messages—both multicast group joins upstream toward a source and data forwarding downstream. Scoping can relieve stress on scarce resources, such as bandwidth, and improve privacy or scaling properties.

IP multicast implementations can achieve some level of scoping by using the time-to-live (TTL) field in the IP header. However, TTL scoping has proven difficult to implement reliably, and the resulting schemes often are complex and difficult to understand.

Administratively scoped IP multicast provides clearer and simpler semantics for multicast scoping. Packets addressed to administratively scoped multicast addresses do not cross configured administrative boundaries. Administratively scoped multicast addresses are locally assigned, and hence are not required to be unique across administrative boundaries.

The administratively scoped IP version 4 (IPv4) multicast address space is the range from 239.0.0.0 through 239.255.255.255.

The structure of the IPv4 administratively scoped multicast space is based loosely on the IP version 6 (IPv6) addressing architecture described in RFC 1884, *IP Version 6 Addressing Architecture*.

There are two well-known scopes:

- IPv4 local scope—This scope comprises addresses in the range 239.255.0.0/16. The local scope is the minimal enclosing scope and is not further divisible. Although the exact extent of a local scope is site-dependent, locally scoped regions must not span any other scope boundary and must be contained completely within or be equal to any larger scope. If scope regions overlap in an area, the area of overlap must be within the local scope.
- IPv4 organization local scope—This scope comprises 239.192.0.0/14. It is the space from which an organization allocates subranges when defining scopes for private use.

The ranges 239.0.0.0/10, 239.64.0.0/10, and 239.128.0.0/10 are unassigned and available for expansion of this space.

Two other scope classes already exist in IPv4 multicast space: the statically assigned link-local scope, which is 224.0.0.0/24, and the static global scope allocations, which contain various addresses.

All scoping is inherently bidirectional in the sense that join messages and data forwarding are controlled in both directions on the scoped interface.

You can configure multicast scoping either by creating a named scope associated with a set of routing device interfaces and an address range, or by referencing a scope policy that specifies the interfaces and configures the address range as a series of filters. You cannot combine the two methods (the commit operation fails for a configuration that includes both). The methods differ somewhat in their requirements and result in different output from the show multicast scope command. For details and configuration instructions, see “Creating a Named Scope for Multicast Scoping” on page 194 and “Using a Scope Policy for Multicast Scoping” on page 195.

For information about supported standards for multicast scoping, see “IP Multicast Specifications” on page 18.

Creating a Named Scope for Multicast Scoping

To configure multicast scoping by defining a named scope, include the scope statement and specify a name for the scope, the set of routing device interfaces on which you are configuring scoping, and the scope's address range:

```
scope scope-name {
    interface [ interface-names ];
    prefix destination-prefix;
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

NOTE: The prefix specified with the prefix statement must be unique for each scope statement. If multiple scopes contain the same prefix, only the last scope applies to the interfaces. If you need to scope the same prefix on multiple interfaces, list all of them in the interface statement for a single scope statement.

When you configure multicast scoping with a named scope, all scope boundaries must include the local scope. If this scope is not configured, it is added automatically at all scoped interfaces. The local scope limits the use of the multicast group 239.255.0.0/16 to an attached LAN.

When you configure scoping with a named scope, the show multicast scope operational mode command displays the names of the defined scopes, prefixes, and interfaces. When you configure scoping with a scope policy, the command displays only the name of the scope policy. For information about scope policies and the scope-policy statement, see "Using a Scope Policy for Multicast Scoping" on page 195.

Using a Scope Policy for Multicast Scoping

To configure multicast scoping with a scope policy, include the scope-policy statement and specify one or more policy names for the scope:

```
scope-policy [ policy-names ];
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Each referenced policy must be correctly configured at the [edit policy-options] hierarchy level, specifying the set of routing device interfaces on which to configure scoping, and defining the scope's address range as a series of route filters. Only the interface, route-filter, and prefix-list match conditions are supported for multicast scope policies. All other configured match conditions are ignored. The only actions supported are accept, reject, and the policy flow actions next-term and next-policy. The reject action means that joins and multicast forwarding are suppressed in both directions on the configured interfaces. The accept action allows joins and multicast forwarding in both directions on the interface. By default, scope policies apply to all interfaces. The default action is accept.

For more information about configuring route filters and policies, see the *JUNOS Policy Framework Configuration Guide*.

NOTE:

Multicast scoping configured with a scope policy differs in several ways from scoping configured with a named scope (which uses the scope statement):

- You cannot apply a scope policy to a specific routing instance, because all scope policies apply to all routing instances. In contrast, a named scope does apply individually to a specific routing instance.
- In contrast to scoping with a named scope, scoping with a scope policy does not automatically add the local scope at scope boundaries. You must explicitly configure the local scope boundaries. The local scope limits the use of the multicast group 239.255.0.0/16 to an attached LAN.
- When you configure multicast scoping with a scope policy, the show multicast scope operational mode command displays only the name of the scope policy. When you configure scoping with a named scope, the command displays the names of the defined scopes, prefixes, and interfaces.

For information about named scopes and the scope statement, see "Creating a Named Scope for Multicast Scoping" on page 194.

Example: Creating a Named Scope for Multicast Scoping

This example configures multicast scoping with the scope statement, creating four scopes: local, organization, engineering, and marketing.

If you have a Tunnel Physical Interface Card (PIC) in your router and you configure a tunnel interface to use IP-IP encapsulation, you can configure the local scope. For more information about configuring tunnel interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

Configure the organization scope on an IP-IP encapsulation tunnel interface and a SONET/SDH interface. Configure the engineering and marketing scopes on an IP-IP encapsulation tunnel interface and two SONET/SDH interfaces. The JUNOS Software can scope any user-configurable IPv6 or IPv4 group.

```
[edit]
routing-options {
   multicast {
      scope local {
         interface gr-2/1/0;
         prefix fe00::239.255.0.0/128;
      }
      scope organization {
         interface [ gr-2/1/0 so-0/0/0 ];
         prefix 239.192.0.0/14;
      }
      scope engineering {
         interface [ ip-2/1/0 so-0/0/1 so-0/0/2 ];
         prefix 239.255.255.0/24;
```

```
        }
        scope marketing {
            interface [ gr-2/1/0 so-0/0/2 so-1/0/0 ];
            prefix 239.255.254.0/24;
        }
    }
}
```

To verify that group scoping is in effect, issue the show multicast scope command:

```
user@host> show multicast scope
Resolve
Scope nameGroup prefixInterfaceRejects
localfe00::239.255.0.0/128gr-2/1/00
organization239.192.0.0/14gr-2/1/0 so-0/0/00
engineering239.255.255.0/24ip-2/1/0 so-0/0/1 so-0/0/20
marketing239.255.254.0/24gr-2/1/0 so-0/0/2 so-1/0/00
```

For more information about the show multicast scope command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Example: Using a Scope Policy for Multicast Scoping

This example configures a scope-policy statement named allow-auto-rp-on-backbone, allowing packets for auto-RP groups 224.0.1.39/32 exact and 224.0.1.40/32 exact on backbone-facing interfaces, and rejecting all other addresses in the 224.0.1.0/24 or longer and 239.0.0.0/8 or longer address ranges.

First, configure the policy allow-auto-rp-on-backbone at the [edit policy-options] hierarchy level:

```
[edit]
policy-options {
    policy-statement allow-auto-rp-on-backbone {
        term allow-auto-rp {
            from {
                /* backbone-facing interfaces */
                interface [ so-0/0/0.0 so-0/0/1.0 ];
                route-filter 224.0.1.39/32 exact;
                route-filter 224.0.1.40/32 exact;
            }
            then {
                accept;
            }
        }
        term reject-these {
            from {
                route-filter 224.0.1.0/24 orlonger;
                route-filter 239.0.0.0/8 orlonger;
            }
            then reject;
        }
    }
}
```

Then apply the scope policy allow-auto-rp-on-backbone at the [edit routing-options] hierarchy level:

```
[edit]
routing-options {
   multicast {
      scope-policy allow-auto-rp-on-backbone;
   }
}
```

To verify that the scope policy is in effect, issue the show multicast scope configuration mode command:

```
user@host> show multicast scope
Scope policy: [ allow-auto-rp-on-backbone ]
```

For more information about the show multicast scope command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Chapter 21

Multicast Reverse Path Forwarding Configuration Guidelines

This chapter discusses the following topics that provide information about configuring multicast RPF policies:

- Multicast RPF Configuration Guidelines on page 199
- Configuring RPF Policies on page 200
- Example: Configuring RPF Policies on page 200

Multicast RPF Configuration Guidelines

You use multicast reverse path forwarding (RPF) checks to prevent multicast routing loops. Routing loops are particularly debilitating in multicast applications because packets are replicated with each pass around the routing loop.

In general, a router is to forward a multicast packet only if it arrives on the interface closest (as defined by a unicast routing protocol) to the origin of the packet, whether source host or rendezvous point (RP). In other words, if a unicast packet would be sent to the "destination" (the reverse path) on the interface that the multicast packet arrived on, the packet passes the RPF check and is processed. Multicast (or unicast) packets that fail the RPF check are not forwarded (this is the default behavior). For an overview of how a Juniper Networks router implements RPF checks with tables, see "Multicast RPF Checks" on page 13.

However, there are network router configurations where multicast packets that fail the RPF check need to be forwarded. For example, when point-to-multipoint label-switched paths (LSPs) are used for distributing multicast traffic to PIM "islands" downstream from the egress router, the interface on which the multicast traffic arrives is not always the RPF interface. This is because LSPs do not follow the normal next-hop rules of independent packet routing. For information on LSPs, see the *JUNOS MPLS Applications Configuration Guide*.

In cases such as these, you can configure policies on the PE router to decide which multicast groups and sources are exempt from the default RPF check.

For more information about policies, see the *JUNOS Policy Framework Configuration Guide*.

Configuring RPF Policies

You configure one or more multicast RPF policies to disable RPF checks for a particular multicast (S,G) pair. You usually disable RPF checks on egress routing devices of a point-to-multipoint LSP.

An RPF policy behaves like an import policy. If no policy term matches the input packet, the default action is to accept (that is, to perform the RPF check).

NOTE: Be careful when disabling RPF checks on multicast traffic. If you disable RPF checks in some configurations, multicast loops can result.

To configure multicast RPF policies, include the rpf-check-policy statement with a correctly configured policy:

```
rpf-check-policy [ policy-names ];
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Changes to an RPF check policy take effect immediately:

- If no policy was previously configured, the policy takes effect immediately.
- If the policy name is changed, the new policy takes effect immediately and any packets no longer filtered are subjected to the RPF check.
- If the policy is deleted, all packets formerly filtered are subjected to the RPF check.
- If the underlying policy is changed, but retains the same name, the new conditions take effect immediately and any packets no longer filtered are subjected to the RPF check.

Example: Configuring RPF Policies

This example configures an RPF check policy named disable-RPF-on-PE, disabling the RPF check on multicast packets for the configured (S,G) source-group pair. This policy will not perform RPF checks on packets arriving for group 228.0.0.0/8 or from source address 196.168.25.6.

First, configure the policy disable-RPF-on-PE at the [edit policy-options] hierarchy level:

```
[edit]
policy-options {
   policy-statement disable-RPF-on-PE {
      term first {
         from {
            route-filter 228.0.0.0/8 orlonger;
            source-address-filter 192.168.25.6/32 exact;
         }
```

```
            then {
                reject;
            }
        }
    }
}
```

For more information about route and source address filters, see the *JUNOS Policy Framework Configuration Guide*.

Then apply the policy disable-RPF-on-PE at the [edit routing-options] hierarchy level:

```
[edit]
routing-options {
    multicast {
        rpf-check-policy disable-RPF-on-PE;
    }
}
```

You can also configure each condition as a separate policy and reference both policies in the rpf-check-policy statement:

```
[edit]
policy-options {
    policy-statement disable-RPF-on-group {
        term first {
            from {
                route-filter 228.0.0.0/8 orlonger;
            }
            then {
                reject;
            }
        }
    }
}
policy-statement disable-RPF-on-source {
    term first {
        from {
            source-address-filter 192.168.25.6/32 exact;
        }
        then {
            reject;
        }
    }
}
[edit]
routing-options {
    multicast {
        rpf-check-policy [ disable-RPF-on-group disable-RPF-on-source ];
    }
}
```

This allows you to associate groups in one policy and sources in the other.

Chapter 22

Source-Specific Multicast Configuration Guidelines

This chapter discusses the following topics:

- Source-Specific Multicast Groups Overview on page 203
- Configuring Source-Specific Multicast Groups on page 204
- Source-Specific Multicast Examples on page 206
- Example: Configuring an SSM-Only Domain on page 206
- Example: Configuring PIM SSM on a Network on page 206
- Example: Configuring SSM Mapping on page 208

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 the JUNOS Software 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 on 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 configure SSM mapping from IGMPv1/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).

For information about standards supported for source-specific multicast, see “IP Multicast Specifications” on page 18.

Configuring Source-Specific Multicast Groups

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including a list of IP addresses with the [edit routing-options multicast ssm-groups] statement.

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 11 on page 204). 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 11 on page 204 that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 11: Receiver Sends Messages to Join Group G and Source S

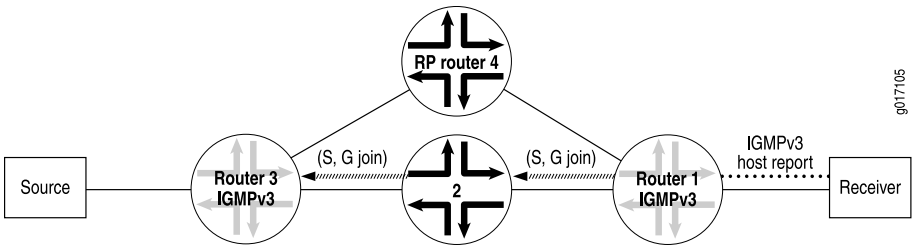

The (S,G) join message initiates the source tree, then builds it out hop by hop until it reaches the source. In Figure 12 on page 205, the source tree is built across the network to Router 3, the last-hop router connected to the source.

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

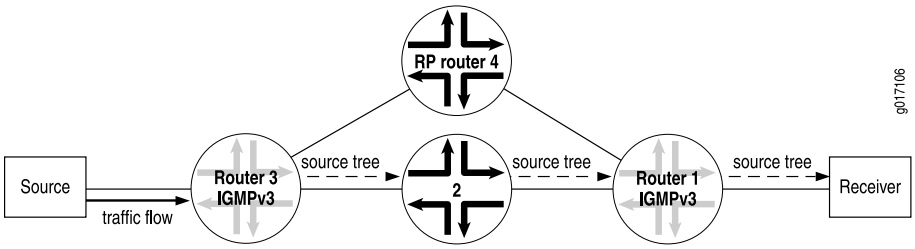

Using the source tree, multicast traffic is delivered to the subscribing host (see Figure 13 on page 205).

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

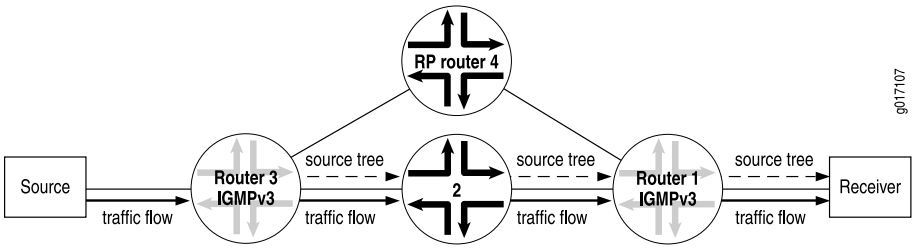

SSM can operate in include mode as previously described or in exclude mode. In exclude mode the receiver specifies a list of sources that it does not want to receive the multicast group traffic from. The routing device forwards traffic to the receiver from any source except the sources specified in the exclusion list. The receiver accepts traffic from any sources except the sources specified in the exclusion list.

To configure additional SSM groups, include the ssm-groups statement:

```
ssm-groups [ 232.0.0.0/8 233.0.0.0/8 239.0.0.0/8 ];
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can configure the JUNOS Software 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.

To configure the JUNOS Software to accept any-source multicast join messages for groups within the SSM address range, include the asm-override-ssm statement:

```
asm-override-ssm;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For an overview of logical systems and a detailed example of logical system configuration, see the logical systems chapter of the *JUNOS Feature Guide*.

Source-Specific Multicast Examples

This section discusses the following topics:

- Example: Configuring an SSM-Only Domain on page 206
- Example: Configuring PIM SSM on a Network on page 206
- Example: Configuring SSM Mapping on page 208

Example: Configuring an SSM-Only Domain

Deploying an SSM-only domain is much simpler than deploying an ASM domain; 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 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 14 on page 207.

Figure 14: Network on Which to Configure PIM SSM

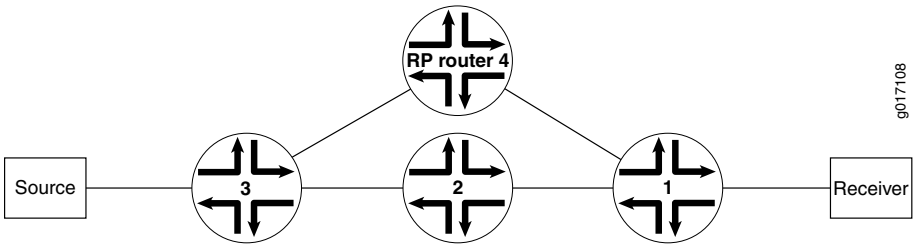

This example configures the IGMP version to IGMPv3 on all receiving host interfaces.

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 and that the fxp0.0 interface has IGMP disabled. (not listed)

   ```
   user@router1> show igmp interface

   Interface      State    Querier      Timeout Version Groups
   fxp1.0         Up                          0       3      0
   fe-0/0/0.0     Up       198.58.3.245     213       3      0
   fe-0/0/1.0     Up       198.58.3.241     220       3      0
   fe-0/0/2.0     Up       198.58.3.237     218       3      0
   Configured Parameters:
   IGMP Query Interval (1/10 secs): 1250
   IGMP Query Response Interval (1/10 secs): 100
   IGMP Last Member Query Interval (1/10 secs): 10
   IGMP Robustness Count: 2
   Derived Parameters:
   IGMP Membership Timeout (1/10 secs): 2600
   IGMP Other Querier Present Timeout (1/10 secs): 2550
   ```

4. Use the show pim join extensive command to verify the PIM join state on Router 2 and Router 3 (the upstream routers).

```
user@router2> show pim join extensive

232.1.1.1       10.4.1.2                        sparse
     Upstream interface: fe-1/1/3.0
     Upstream State: Local Source
     Keepalive timeout: 209
     Downstream Neighbors:
        Interface: so-1/0/2.0
            10.10.71.1       State: Join   Flags: S    Timeout: 209
```

5. Use the show pim join extensive command to verify the PIM join state on Router 1 (the router connected to the receiver).

```
user@router1> show pim join extensive

232.1.1.1       10.4.1.2                        sparse
     Upstream interface: so-1/0/2.0
     Upstream State: Join to Source
     Keepalive timeout: 209
     Downstream Neighbors:
        Interface: fe-0/2/3.0
            10.3.1.1         State: Join   Flags: S    Timeout: Infinity
```

Example: Configuring SSM Mapping

SSM mapping does not require all hosts to support IGMPv3. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report. This allows hosts running IGMPv1/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 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. Creates 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
```

Chapter 23

Flow Map Configuration Guidelines

This chapter contains the following sections:

- Creating a Multicast Flow Map on page 211
- Defining Flow Properties on page 212
- Displaying Flow Maps on page 213

Creating a Multicast Flow Map

Multicast flow maps enable you to manage a subset of multicast forwarding table entries. For example, you can specify that certain forwarding cache entries be permanent or have a different timeout value from other multicast flows that are not associated with the flow map policy.

Creating a Flow Map Policy

The following example creates a flow map policy called policyForFlow1. In this example, the policy statement matches the source address using the source-address-filter statement and matches the group address using the prefix-list-filter.

```
[edit policy-options]
prefix-list permanentEntries1 {
   232.1.1.0/24;
}
policy policyForFlow1 {
   from {
      source-address-filter 11.11.11.11/32 exact;
      prefix-list-filter permanentEntries1 orlonger;
   }
   then accept;
}
```

NOTE: The addresses must match the configured policy for flow mapping to occur.

For additional information about creating policy statements, see the *JUNOS Policy Framework Configuration Guide*.

Defining Flow Maps

This example defines a flow map called flowMap1 with permanent forwarding entries (that is, the entries never time out).

```
[edit routing-options]
multicast {
   flow-map flowMap1 {
      policy policyForFlow1;
      forwarding-cache {
         timeout never;
      }
   }
}
```

Defining Flow Properties

Flow maps enable you to configure bandwidth variables and multicast forwarding cache timeout values for entries defined by the flow map policy. This section discusses the following flow properties:

- Defining Bandwidth for Multicast Flows on page 212
- Defining Forwarding Cache Timeout on page 212
- Specifying Redundant Flow Sources on page 212

Defining Bandwidth for Multicast Flows

You can define the bandwidth associated with multicast flows that match a flow map by specifying a bandwidth in bits per second, by specifying that the bandwidth is measured and adaptively modified, or by a combination of both methods where you specify an initial value for the bandwidth and the interface uses the adaptive value thereafter. For additional information about how to define flow bandwidth, see "Defining Bandwidth for Multicast Flows" on page 218.

Defining Forwarding Cache Timeout

The forwarding cache timeout value can range from 1 through 720 minutes, or you can set the value to never, making the entries permanent. For additional information about how to configure forwarding cache properties, see "Configuring General Multicast Forwarding Cache Properties" on page 229.

Specifying Redundant Flow Sources

To specify redundant (backup) sources for flows identified by a flow map, include the redundant-sources statement:

```
redundant-sources [ addresses ];
```

You can include this statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*]
- [edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*]
- [edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*]
- [edit routing-options multicast flow-map *flow-map-name*]

Outbound interfaces that are admitted for one of the forwarding entries are automatically admitted for any other entries identified by the redundant source configuration.

In the following example, forwarding entries (10.11.11.11, g1) and (10.11.11.12, g1) match the flow map flowMap1. In this case, if a particular outbound interface is admitted for entry (10.11.11.11, g1), it is automatically admitted for entry (10.11.11.12, g1), even if there is no longer enough remaining bandwidth available after creating entry (10.11.11.11, g1). The interface is added because only one of the two sources can send traffic at any time.

```
[edit routing-options]
multicast {
   flow-map flowMap1 {
      redundant-sources [ 10.11.11.11 10.11.11.12 ];
   }
}
```

Displaying Flow Maps

Purpose To display configured flow maps, the policies that they use, and their configuration settings, use the show multicast flow-map command:

Action

```
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
```

For additional information about the show multicast flow-map command, including descriptions of the fields, see the *JUNOS Routing Protocols and Policies Command Reference*.

Chapter 24

Bandwidth Management Configuration Guidelines

This chapter discusses the following topics:

- Bandwidth Management Overview on page 215
- Bandwidth Management and PIM Graceful Restart on page 216
- Bandwidth Management and Source Redundancy on page 216
- Logical Systems and Bandwidth Oversubscription on page 217
- How Interface Bandwidth Becomes Oversubscribed on page 217
- How Interface Bandwidth Becomes Available Again on page 217
- Readmitting or Removing Interfaces on page 217
- Defining Interface Bandwidth Maximums on page 218
- Defining Bandwidth for Multicast Flows on page 218
- Examples: Defining Bandwidths on page 219
- Example: Configuring Maximum Multicast Bandwidth on an Interface on page 219
- Example: Configuring Bandwidth for Individual Flows on page 220
- Configuring Multicast Outgoing Interface Mapping on page 220
- Configuring Multicast Routing Over IP Demux Interfaces on page 225
- Managing Subscriber Overcommitment on page 226
- Configuring Multicast Egress Class of Service on page 227

Bandwidth Management Overview

Bandwidth management enables you to control the multicast flows that leave a multicast interface. This control enables you to better manage your multicast traffic and reduce or eliminate the chances of interface oversubscription or congestion.

Bandwidth management ensures that multicast traffic oversubscription does not occur on an interface. When managing multicast bandwidth, you define the maximum amount of multicast bandwidth that an individual interface can use as well as the bandwidth individual multicast flows use.

For example, the routing software cannot add a flow to an interface if doing so exceeds the allowed bandwidth for that interface. Under these circumstances, the

interface is rejected. This rejection, however, does not prevent a multicast protocol (for example, PIM) from sending a join message upstream. Traffic continues to arrive on the router, even though the router is not sending the flow from the expected outgoing interfaces.

You can configure the flow bandwidth statically by specifying a bandwidth value for the flow in bits per second or you can enable the flow bandwidth to be measured and adaptively changed. When using the adaptive bandwidth option, the routing software queries the statistics for the flows to be measured at five-second intervals and calculates the bandwidth based on the queries. The routing software uses the maximum value measured within the last minute (that is, the last twelve measuring points) as the flow bandwidth.

For more information, see the following sections:

- Bandwidth Management and PIM Graceful Restart on page 216
- Bandwidth Management and Source Redundancy on page 216
- Logical Systems and Bandwidth Oversubscription on page 217

Bandwidth Management and PIM Graceful Restart

When using PIM graceful restart, after the routing process restarts on the Routing Engine, previously admitted interfaces are always readmitted and the available bandwidth is adjusted on the interfaces. When using the adaptive bandwidth option, the bandwidth measurement is initially based on the configured or default starting bandwidth, which might be inaccurate during the first minute. This means that new flows may be incorrectly rejected or admitted temporarily. You can correct this problem by issuing the clear multicast bandwidth-admission operational command.

If PIM graceful restart is not configured, after the routing process restarts, previously admitted or rejected interfaces might be rejected or admitted in an unpredictable manner.

For additional information about the clear multicast bandwidth-admission command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Bandwidth Management and Source Redundancy

When using source redundancy, multiple sources (for example, s1 and s2) might exist for the same destination group (g). However, only one of the sources can actively transmit at any time. In this case, multiple forwarding entries—(s1,g) and (s2,g)—are created after each goes through the admission process.

With redundant sources, unlike unrelated entries, an OIF that is already admitted for one entry—for example, (s1,g)—is automatically admitted for other redundancy entries—for example, (s2,g). The remaining bandwidth on the interface is deducted each time an outbound interface is added, even though only one sender actively transmits. By measuring bandwidth, the bandwidth deducted for the inactive entries is credited back when the router detects no traffic is being transmitted.

For more information about defining redundant sources, see “Specifying Redundant Flow Sources” on page 212.

Logical Systems and Bandwidth Oversubscription

You can manage bandwidth at both the physical and logical interface level. However, if more than one logical system shares the same physical interface, the interface might become oversubscribed. Oversubscription occurs if the total bandwidth of all separately configured maximum bandwidth values for the interfaces on each logical system exceeds the bandwidth of the physical interface.

When displaying interface bandwidth information, a negative available bandwidth value indicates oversubscription on the interface.

This section discusses the following logical system oversubscription topics:

- How Interface Bandwidth Becomes Oversubscribed on page 217
- How Interface Bandwidth Becomes Available Again on page 217
- Readmitting or Removing Interfaces on page 217

How Interface Bandwidth Becomes Oversubscribed

Interface bandwidth can become oversubscribed under the following conditions:

- The configured maximum bandwidth decreases.
- Some flow bandwidths increase because of a configuration change or an actual increase in the traffic rate.

How Interface Bandwidth Becomes Available Again

Interface bandwidth can become available again under the following conditions:

- The configured maximum bandwidth increases.
- Some flows are no longer transmitted from interfaces and bandwidth reserves for them are now available to other flows.
- Some flow bandwidths decrease because of a configuration change or an actual decrease in the traffic rate.

Readmitting or Removing Interfaces

Interfaces rejected for a flow due to insufficient bandwidth are not automatically readmitted, even when bandwidth becomes available again. Rejected interfaces have an opportunity to be readmitted when one of the following occurs:

- The multicast routing protocol updates the forwarding entry for the flow after receiving a join, leave, or prune message or a topology change occurs.

- The multicast routing protocol updates the forwarding entry for the flow due to configuration changes.
- You manually reapply bandwidth management to a specific flow or to all flows using the clear multicast bandwidth-admission operational command.

In addition, even if previously available bandwidth is no longer available, already admitted interfaces are not removed until one of the following occurs:

- The multicast routing protocol explicitly removes the interfaces after receiving a leave or prune message or a topology change occurs.
- You manually reapply bandwidth management to a specific flow or to all flows using the clear multicast bandwidth-admission operational command.

For additional information about the clear multicast bandwidth-admission command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Defining Interface Bandwidth Maximums

To define the maximum multicast bandwidth for an interface, include the maximum-bandwidth statement:

maximum-bandwidth *bps*;

NOTE: You need only define the maximum bandwidth for an interface on which you want to apply bandwidth management. An interface that does not have a defined maximum bandwidth transmits all multicast flows as determined by the multicast protocol that is running on the interface (for example, PIM).

You can include this statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*]
- [edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*]
- [edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*]
- [edit routing-options multicast interface *interface-name*]

Defining Bandwidth for Multicast Flows

To specify bandwidth for each flow that is defined by a flow map, include the bandwidth statement. You can define the bandwidth associated with a multicast flow statically by specifying a bandwidth in bits per second or by specifying that the routing software determine bandwidth using adaptive bandwidth measuring. To specify the use of adaptive bandwidth measuring, use the adaptive option.

bandwidth (*bps* | adaptive);

You can include this statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*]
- [edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*]
- [edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*]
- [edit routing-options multicast flow-map *flow-map-name*]

When you use the adaptive option, the bandwidth adjusts based on measurements made at 5-second intervals. The flow uses the maximum bandwidth value from the last twelve measured values (one minute).

You can configure a bandwidth value (in bits per second) together with the adaptive option. In this case, the bandwidth value acts as the starting bandwidth for the flow. The bandwidth then changes based on subsequent measured bandwidth values. If you do not specify a bandwidth value with the adaptive option, the starting bandwidth defaults to 2 megabits per second (Mbps).

For example, the bandwidth 2m adaptive statement is equivalent to the bandwidth adaptive statement because they both use the same starting bandwidth (2 Mbps, the default). If the actual flow bandwidth is 4 Mbps, the measured flow bandwidth changes to 4 Mbps after reaching the first measuring point (5 seconds). However, if the actual flow bandwidth rate is 1 Mbps, the measured flow bandwidth remains at 2 Mbps for the first twelve measurement cycles (one minute) and then changes to the measured 1 Mbps value.

Examples: Defining Bandwidths

For examples of how to manage multicast bandwidth, see the following sections:

- Example: Configuring Maximum Multicast Bandwidth on an Interface on page 219
- Example: Configuring Bandwidth for Individual Flows on page 220

Example: Configuring Maximum Multicast Bandwidth on an Interface

Configure three multicast interface bandwidth maximums. The maximum bandwidth of the first logical interface is the link speed (100 Mbps), the maximum bandwidth of the second physical interface is 60 Mbps, and the maximum bandwidth of the third logical interface is 10 Mbps:

```
[edit]
   routing-options {
   multicast {
      interface fe-0/2/0.200 {
         maximum-bandwidth;
      }
      interface fe-0/2/1 {
```

```
            maximum-bandwidth 60m;
        }
        interface fe-0/2/1.210 {
            maximum-bandwidth 10m;
        }
    }
}
```

Example: Configuring Bandwidth for Individual Flows

Configure the flow map bandwidth to be adaptive with a default starting bandwidth of 2 Mbps:

```
[edit]
    routing-options {
    multicast {
        flow-map map1 {
            policy policy1;
            bandwidth 2m adaptive;
        }
    }
}
```

Configuring Multicast Outgoing Interface Mapping

Multicast outgoing interface (OIF) mapping enables you to map multicast output to different interfaces. You can use this feature to enable the router to decrease the inefficiencies associated with replicating streams of multicast traffic by directing multicast streams (data packets) to a different interface than the one used for multicast control packets (used for IGMP and MLD joins and leaves). If multiple control interfaces all request the same stream and are mapped to the same output interface, only one copy of the stream is transmitted. Once an OIF map has been configured, you can apply it to an IGMP interface, enabling it to use a different output interface for multicast data forwarding.

NOTE: OIF mapping and reverse OIF mapping are not supported on the same customer or shared interface.

Configuring OIF mapping on a multicast interface requires the following general steps:

1. Configuring an OIF Map on page 221
2. Associating OIF Maps with Interfaces on page 222
3. Disabling QoS Adjustment on page 223
4. Configuring Passive Mode for IGMP or MLD on page 224

Configuring an OIF Map

The OIF map is a routing policy statement that can contain multiple terms. When creating OIF maps, keep the following in mind:

- If you specify a physical interface as the map-to-interface statement value (for example, ge-0/0/0), a ".0" is appended to the interface to create a logical interface (for example, ge-0/0/0.0).
- Configure a routing policy for each logical system. You cannot configure routing policies dynamically. As part of the policy, you must configure the map-to-interface statement statically.
- To use OIF mapping, you must also have IGMP, MLD, or PIM configured.
- You cannot map to a mapped interface.
- To use call admission control (CAC) with a mapped interface, you must configure it using the maximum-bandwidth statement (see the existing documentation for details about this statement).
- We recommend that you configure policy statements for IGMP and MLD separately.

To specify an OIF map:

1. Access the [edit policy-options] hierarchy.

   [edit]
   user@host# **edit policy-options**

2. Specify the policy that you want to use as an OIF map.

   [edit policy-options]
   user@host# **edit policy-statement map1**

3. Include any desired terms and specify any from or then clauses within them. The from clause in each term can select multicast flows based on the multicast group (using the route-filter keyword) and the source address (with the source-address-filter keyword). The then clause must accept or reject the term.

   [edit policy-options policy-statement map1]
   user@host# **set from family inet**
   user@host# **set then accept**

NOTE: For detailed information about how to create policies, see the *JUNOS Policy Framework Configuration Guide*.

---

4. Access the policy statement term.

   [edit policy-options policy-statement map1]
   user@host# **edit term term1**

5. Include a then action for the map-to-interface statement. The map-to-interface statement sets a value similar to the existing metric or tag actions and requires that you specify either a logical interface (that is, any interface that multicast currently supports) or the keyword self.

   [edit policy-options policy-statement map1 term term1]
   user@host# **edit then map-to-interface**

6. Specify either a logical interface (that is, any logical interface that multicast currently supports) or the keyword self. The self keyword specifies that multicast data packets be sent on the same interface as the control packets and that no mapping occur. If no term matches, then no multicast data packets are sent.

   [edit policy-options policy-statement map1 term term1 then map-to-interface]
   user@host# **set self**

NOTE: For additional information about how to configure the map-to-interface statement, see "Configuring Actions in Routing Policy Terms" in the *JUNOS Policy Framework Configuration Guide*.

Associating OIF Maps with Interfaces

For OIF mapping to function, you must associate the configured OIF map with a logical interface. The OIF map is then applied to all IGMP or MLD requests received on the configured interface.

NOTE: If an OIF map is already configured on an interface, the new OIF map is added to the policy.

To associate an OIF map with a logical interface:

1. Access the IGMP or MLD interface.

   [edit]
   user@host# **edit protocols igmp interface ge-1/0/0.1**

   or

   [edit]
   user@host# **edit protocols mld interface ge-1/0/0.2**

2. Configure the OIF map that you want to associate with the interface as follows:
   - For IGMP interfaces, include the IGMP oif-map statement.
   - For MLD interfaces, include the MLD oif-map statement.

   [edit protocols igmp interface ge-1/0/0.1]
   user@host# **set oif-map MAP1**

or

```
[edit protocols mld interface ge-2/0/0.1]
user@host# set oif-map MAP2
```

NOTE: You can dynamically associate OIF maps with IGMP interfaces using dynamic profiles. For detailed information about creating dynamic profiles, see the *JUNOS Subscriber Access Configuration Guide*.

Disabling QoS Adjustment

QoS adjustment decreases the available bandwidth on the client interface by the amount of bandwidth consumed by the multicast streams that are mapped from the client interface to the shared interface. This action always occurs unless it is explicitly disabled.

You can optionally disable QoS adjustment. If you do so, available bandwidth is not reduced on the customer interface when multicast streams are added to the shared interface.

To disable QoS adjustment:

1. Access the multicast interface hierarchy.

   ```
   [edit]
   user@host# edit routing-options multicast interface ge-1/0/1
   ```

2. Include the no-qos-adjust statement to disable QoS adjustment.

   ```
   [edit routing-options multicast interface ge-1/0/1]
   user@host# set no-qos-adjust
   ```

NOTE: You can dynamically disable QoS adjustment for IGMP interfaces using dynamic profiles. For more detailed information about creating dynamic profiles, see the *JUNOS Subscriber Access Configuration Guide*.

Configuring Passive Mode for IGMP or MLD

You can optionally define either IGMP or MLD to use passive mode.

NOTE: The OIF map interface should not typically pass IGMP or MLD control traffic and should be configured as passive. However, the OIF map implementation does support running IGMP or MLD on an interface (control and data) in addition to mapping data streams to the same interface. In this case, you should configure IGMP or MLD normally (that is, not in passive mode) on the mapped interface.

By default, specifying the passive statement means that no general queries, group-specific queries, or group-source-specific queries are sent over the interface and that all received control traffic is ignored by the interface. However, you can selectively activate up to two out of the three available options for the passive statement while keeping the other functions passive (inactive). These options include the following:

- send-general-query—When specified, the interface sends general queries.
- send-group-query—When specified, the interface sends group-specific and group-source-specific queries.
- allow-receive—When specified, the interface receives control traffic.

To configure passive mode for IGMP or MLD interfaces:

1. Access the IGMP or MLD interface.

   ```
   [edit]
   user@host# edit protocols igmp interface ge-1/0/0.1
   ```

   or

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

2. Configure the interface to use passive mode as follows:
   - For IGMP interfaces, include the IGMP passive statement.
   - For MLD interfaces, include the MLD passive statement.

   ```
   [edit protocols igmp interface ge-1/0/0.1]
   user@host# set passive
   ```

   or

   ```
   [edit protocols mld interface ge-2/0/0.1]
   user@host# set passive
   ```

NOTE: You can dynamically define IGMP interfaces as passive. For more detailed information about creating dynamic profiles, see the *JUNOS Subscriber Access Configuration Guide*.

Related Topics

- Configuring Multicast Routing Over IP Demux Interfaces on page 225
- *JUNOS Subscriber Access Configuration Guide*

Configuring Multicast Routing Over IP Demux Interfaces

In a subscriber management network, fields in packets sent from IP demux interfaces are intended to correspond to a specific client that resides on the other side of an aggregation device (for example, a Multiservice Access Node [MSAN]). However, packets sent from a Broadband Services Router (BSR) to an MSAN do not identify the demux interface. Once it obtains a packet, it is up to the MSAN device to determine which client receives the packet.

Depending on the intelligence of the MSAN device, determining which client receives the packet can occur in an inefficient manner. For example, when it receives IGMP control traffic, an MSAN might forward the control traffic to all clients instead of the one intended client. In addition, once a data stream destination is established, though an MSAN can use IGMP snooping to determine which hosts reside in a particular group and limit data streams to only that group, the MSAN still must send multiple copies of the data stream to each group member, even if that data stream is intended for only one client in the group.

Various multicast features, when combined, enable you to avoid the inefficiencies mentioned above. These features include the following:

- The ability to configure the IP demux interface family statement to use inet for either the numbered or unnumbered primary interface. See Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles for details.
- The ability to configure IGMP on the primary interface to send general queries for all clients. The demux configuration prevents the primary IGMP interface from receiving any client IGMP control packets. Instead, all IGMP control packets go to the demux interfaces. However, to guarantee that no joins occur on the primary interface:
  - For static IGMP interfaces—Include the passive send-general-query statement in the IGMP configuration at the [edit protocols igmp interface *interface-name*] hierarchy level.
  - For dynamic IGMP demux interfaces—Include the passive send-general-query statement at the [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*] hierarchy level.
- The ability to map all multicast groups to the primary interface as follows:
  - For static IGMP interfaces—Include the oif-map statement at the [edit protocols igmp interface *interface-name*] hierarchy level.

- For dynamic IGMP demux interfaces—Include the oif-map statement at the [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*] hierarchy level.

Using the oif-map statement, you can map the same IGMP group to the same output interface and send only one copy of the multicast stream from the interface.

- The ability to configure IGMP on each demux interface. To prevent duplicate general queries:
  - For static IGMP interfaces—Include the passive allow-receive send-group-query statement at the [edit protocols igmp interface *interface-name*] hierarchy level.
  - For dynamic demux interfaces—Include the passive allow-receive send-group-query statement at the [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*] hierarchy level.

NOTE: To send only one copy of each group, regardless of how many customers join, use the oif-map statement as previously mentioned.

Related Topics

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles
- Configuring Multicast Outgoing Interface Mapping on page 220
- *JUNOS Subscriber Access Configuration Guide*

Managing Subscriber Overcommitment

The reverse-oif-mapping statement enables the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD join or leave request it receives over the multicast VLAN. Once the subscriber VLAN is identified, the routing device immediately adjusts the quality of service (in this case, the bandwidth) on that VLAN based on the addition or removal of a subscriber.

By default, hierarchical bandwidth adjustment is performed for all subscriber interfaces with reverse OIF mapping enabled on the specified multicast interface. You can disable hierarchical bandwidth adjustment by including the no-qos-adjust statement.

To introduce a delay to the QoS update, use the subscriber-leave-timer statement. This statement defines a time delay (between 1 and 30 seconds) that the routing device waits before updating the QoS for the remaining subscriber interfaces after receiving an IGMP or MLD leave request. You might use this delay to decrease how often the routing device adjusts the overall QoS bandwidth on the VLAN when a subscriber sends rapid leave and join messages (for example, when changing channels in an IPTV network).

NOTE: The routing device uses IGMP and MLD join or leave reports to obtain the subscriber VLAN information. This means that the connecting equipment (for example, the DSLAM) must forward all IGMP and MLD reports to the routing device in order for this feature to function properly. Using report suppression or an IGMP proxy can result in reverse OIF mapping not working properly.

```
interface interface-name {
    reverse-oif-mapping {
        no-qos-adjust;
    }
    subscriber-leave-timer seconds;
}
```

You can include this statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast]
- [edit logical-systems *logical-system-name* routing-options multicast]
- [edit routing-instances *routing-instance-name* routing-options multicast]
- [edit routing-options multicast]

Related Topics

- Hierarchical CoS Shaping-Rate Adjustments Overview in the JUNOS® Software Subscriber Access Configuration Guide

Configuring Multicast Egress Class of Service

For the IQ, IQE, LSQ, and ATM2 PICs in M320, T320, and T640 routers, you can configure a forwarding class map and queue number for unicast and multicast traffic on an egress interface.

To configure a interface-specific forwarding-class map, include the forwarding-classes-interface-specific statement at the [edit class-of-service] hierarchy level.

To apply the forwarding-class map to a logical interface, include the output-forwarding-class-map forwarding-class-map-name statement at the [edit interfaces interface-name unit logical-unit-number] hierarchy level.

For PICs that are restricted to four queues, you can control the queue assignment with the restricted-queues statement or rely on the routing platform to override the queue assignment using modular arithmetic.

For more information about configuring multicast forwarding class maps, see the *Class of Service Configuration Guide*.

Chapter 25

Multicast Forwarding Cache Configuration Guidelines

This section discusses the following topics, which provide information about configuring multicast forwarding cache properties:

- Multicast Forwarding Cache Overview on page 229
- Configuring General Multicast Forwarding Cache Properties on page 229
- Configuring Multicast Forwarding Cache Properties for Flow Maps on page 230
- Examples: Configuring Multicast Forwarding Cache Properties on page 231
- Configuring Forwarding Cache Properties at the Multicast Level on page 231
- Configuring Forwarding Cache Properties at the Flow Map Level on page 232
- Displaying the Cache Timeout on page 233

Multicast Forwarding Cache Overview

IP multicast protocols can create numerous entries in the multicast forwarding cache. If the forwarding cache fills up with entries that prevent the addition of higher-priority entries, applications and protocols might not function properly. You can manage the multicast forwarding cache properties by limiting the size of the cache and by controlling the length of time that entries remain in the cache. By managing timeout values, you can give preference to more important forwarding cache entries while removing other less important entries.

Configuring General Multicast Forwarding Cache Properties

To configure multicast forwarding cache limits and timeout values, include the forwarding-cache statement:

```
forwarding-cache {
    threshold suppress value <reuse value>;
    timeout minutes;
}
```

You can include this statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast]
- [edit logical-systems *logical-system-name* routing-options multicast]
- [edit routing-instances *routing-instance-name* routing-options multicast]
- [edit routing-options multicast]

For an overview of logical systems and a detailed example of logical system configuration, see the logical systems chapter of the *JUNOS Feature Guide*.

By default, there are no limits on the number of multicast forwarding cache entries.

You can specify a value for the threshold to suppress new multicast forwarding cache entries and an optional reuse value for the threshold at which the routing device begins to create new multicast forwarding cache entries. The range for both is from 1 through 200,000. If configured, the reuse value must be less than the suppression threshold value. The suppression value is mandatory. If you do not specify the optional reuse value, then the number of multicast forwarding cache entries is limited to the suppression value. A new entry is created as soon as the number of multicast forwarding cache entries falls below the suppression value.

You can also specify a timeout value for all multicast forwarding cache entries. The range for the timeout value is 1 through 720 minutes.

For information about supported standards for multicast scoping, see "IP Multicast Specifications" on page 18.

Configuring Multicast Forwarding Cache Properties for Flow Maps

The forwarding-cache statement enables you to configure the forwarding cache properties of entries defined by a flow map. You can specify a timeout of never to make the forwarding entries permanent, or you can specify a timeout in the range from 1 through 720 minutes.

You can use longer timeouts or permanent forwarding cache entries in source redundancy scenarios to decrease the time delay inherent in switching from one source to another. With source redundancy, only one source (for example, s1) actively sends traffic at any given time, but the sources can switch. This means that the forwarding state for another source (for example, s2) can time out.

Even though the Routing Engine has the corresponding PIM states set up, when the first (s2,g) packet arrives on each routing device after a switchover from source s1 to source s2, the routing device must reinstall the forwarding path (s2,g) if the path has timed out. If many data streams exist, the switchover can take considerable time. Using longer timeouts or permanent cache entries helps reduce the inherent switchover delay.

To configure the multicast forwarding cache timeout for a specified flow map, include the forwarding-cache statement:

```
forwarding-cache {
   timeout (minutes | never );
```

```
}
```

You can include this statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*]
- [edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*]
- [edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*]
- [edit routing-options multicast flow-map *flow-map-name*]

For an overview of logical systems and a detailed example of logical system configuration, see the logical systems chapter of the *JUNOS Feature Guide*.

NOTE: The permanent forwarding state must exist on all routing devices in the path for fast source switchover to function properly.

For information about configuring multicast flow maps, see "Creating a Multicast Flow Map" on page 211.

Examples: Configuring Multicast Forwarding Cache Properties

Depending on the hierarchy level at which you modify the multicast forwarding cache properties, you can specify a threshold value, a timeout value, or both. See the following sections:

- Configuring Forwarding Cache Properties at the Multicast Level on page 231
- Configuring Forwarding Cache Properties at the Flow Map Level on page 232

Configuring Forwarding Cache Properties at the Multicast Level

Configure multicast forwarding cache limits, establishing a suppression threshold of 10, a reuse threshold of 7, and no timeout value:

```
routing-options {
   multicast {
      forwarding-cache {
         threshold suppress 10 reuse 7;
      }
   }
}
```

Only 10 multicast forwarding cache entries are created. No new multicast forwarding cache entries are created until the number of multicast forwarding cache entries falls to 7.

Configure multicast forwarding cache limits, establishing a suppression threshold of 20, no reuse threshold, and no timeout value:

```
routing-options {
   multicast {
      forwarding-cache {
         threshold suppress 20;
      }
   }
}
```

Only 20 multicast forwarding cache entries are created. A new multicast forwarding cache entry is created when the number of multicast forwarding cache entries falls to 19.

Configure the multicast forwarding cache, establishing a timeout of 60 minutes, no suppression, and no reuse threshold:

```
routing-options {
   multicast {
      forwarding-cache {
         timeout 60;
      }
   }
}
```

Multicast forwarding cache entries are created as necessary. Forwarding cache entries are deleted after being idle for 60 minutes (1 hour).

Configuring Forwarding Cache Properties at the Flow Map Level

Configure the flow map forwarding cache, establishing a timeout of 120 minutes:

```
routing-options {
   multicast {
      flow-map flowMap1
      forwarding-cache {
         timeout 120;
      }
   }
}
```

Multicast forwarding cache entries associated with flow map flowMap1 are deleted after being idle for 120 minutes (2 hours).

Configure the flow map forwarding cache, establishing permanent forwarding cache entries:

```
routing-options {
   multicast {
      flow-map flowMap2
      forwarding-cache {
         timeout never;
      }
```

```
    }
  }
```

Multicast forwarding cache entries associated with flow map flowMap2 are permanent (that is, they never time out).

Displaying the Cache Timeout

Purpose To display the cache timeout value for each multicast route, use the show multicast route extensive command:

Action

```
user@host> show multicast route extensive
Family: INET
Group: 232.0.0.1
    Source: 11.11.11.11/32
    Upstream interface: fe-0/2/0.200
    Downstream interface list:
        fe-0/2/1.210
    Downstream interface list rejected by CAC:
        fe-0/2/1.220
    Session description: Source specific multicast
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 337
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: never
Wrong incoming interface notifications: 0
```

For additional information about this show command, including descriptions of the fields, see the *JUNOS Routing Protocols and Policies Command Reference*.

Chapter 26

Ingress PE Redundancy Configuration Guidelines

This section discusses the following topics, which provide information about configuring ingress PE redundancy properties:

- Ingress PE Redundancy Configuration Overview on page 235
- Configuring Ingress PE Redundancy on page 235
- Example: Ingress PE Redundancy on page 237

Ingress PE Redundancy Configuration Overview

In many network topologies, point-to-multipoint label-switched paths (LSPs) are used to distribute multicast traffic over a virtual private network (VPN). When traffic engineering is added to the provider edge (PE) routers, a popular deployment option has been to use traffic-engineered point-to-multipoint LSPs at the origin PE. In these network deployments, the PE is a single point of failure. Network operators have previously provided redundancy by broadcasting duplicate streams of multicast traffic from multiple PEs, a practice which at least doubles the bandwidth required for each stream.

Ingress PE redundancy eliminates the bandwidth duplication requirement by configuring one or more ingress PEs as a group. Within a group, one PE is designated as the primary PE and one or more others become backup PEs for the configured traffic stream. The solution depends on a full mesh of point-to-point (P2P) LSPs among the primary and backup PEs. Also, you must configure a full set of point-to-multipoint LSPs at the backup PEs, even though these point-to-multipoint LSPs at the backup PEs are not sending any traffic or using any bandwidth. The P2P LSPs are configured with bidirectional forwarding detection (BFD). When BFD detects a failure on the primary PE, a new designated forwarder is elected for the stream.

For more information about traffic engineering, LSPs, and BFD, see the *JUNOS MPLS Applications Configuration Guide*.

Configuring Ingress PE Redundancy

You configure one or more PEs as part of a backup PE group to enable ingress PE redundancy. You do this by configuring the IP addresses of the backup PEs (at least one backup PE is required) and the local IP address used by the local PE.

NOTE: You must also configure a full mesh of P2P LSPs between the primary and backup PEs. You must also configure these LSPs with BFD. These configuration steps are not discussed in this section. For more information about mesh LSPs and BFD configuration, see the *JUNOS MPLS Applications Configuration Guide*.

To configure ingress PE redundancy for multicast traffic streams, include the backup-pe-group statement:

```
backup-pe-group group-name {
    backups [ addresses ];
    local-address address;
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

After you configure the ingress PE redundancy backup group, you must also apply the group to a static route on the PE. This makes sure that the static route is active (installed in the forwarding table) when the local PE is the designated forwarder for the configure backup PE group. You can only associate a backup PE group with a static route that has the p2mp-lsp-next-hop statement configured. For more information about the p2mp-lsp-next-hop statement, see the *JUNOS MPLS Applications Configuration Guide*.

You must also configure a full mesh of P2P LSPs between the PEs in the backup group. A full mesh is required so that each member of the group can make an independent decision about the health of the other PEs and determine the designated forwarder for the group. This configuration also requires a BFD configuration for the LSP and the associate-backup-pe-groups statement. An LSP with this statement set monitors the health of the routing device at the other end of the LSP. You can configure multiple backup PE groups that contain the same routing device's address. Failure of this LSP indicates to all of these groups that the destination PE routing device is down. So the associate-backup-pe-groups statement is not tied to any specific group but applies to all groups that are monitoring the health of the LSP to the remote address.

If there are multiple LSPs with the associate-backup-pe-groups statement to the same destination PE, then the local routing device picks the first LSP to that PE for detection purposes.

NOTE: We do not recommend configuring multiple LSPs to the same destination. If you do, make sure that the LSP parameters (for example, liveliness detection) are similar in order to avoid false failure notification even when the remote PE is up.

For more information about the associate-backup-pe-groups statement, see the *JUNOS MPLS Applications Configuration Guide*.

For a complete example of ingress PE redundancy configuration, see “Example: Ingress PE Redundancy” on page 237.

Example: Ingress PE Redundancy

This example establishes a backup PE group for multicast ingress PE redundancy. It also includes examples of the **p2mp-lsp-next-hop** and **backup-pe-group** statements for the static route (both are required), as well as the **associate-backup-pe-groups** statement and BFD for the LSP (both are required). This example does not detail the configuration of the LSP mesh among the PE routers.

Configure the backup PE group:

```
[edit routing-options multicast]
backup-pe-group GroupOne {
   backups [ 10.10.10.2 10.10.10.3 ];
   local-address 10.10.10.1;
}
```

Configure the static route for the point-to-multipoint LSP next hop and backup PE group:

```
[edit routing-options]
static {
   route 10.1.1.0/24 {
      p2mp-lsp-next-hop p2mp-lsp-example;
      backup-pe-group GroupOne;
   }
}
```

Configure BFD and the backup PE group for the LSP mesh:

```
[edit protocols mpls]
label-switched-path IngressPEexample {
   to 10.255.165.9;
   oam {
      bfd-liveness-detection {
         minimum-interval 500;
      }
   }
   associate-backup-pe-groups;
}
```

Chapter 27

Configuring PIM-to-IGMP and PIM-to-MLD Message Translation

This chapter discusses the following topics, which provide information about using PIM-to-IGMP message translation and PIM-to-MLD message translation to forward multicast traffic across PIM domains in certain network topologies:

- PIM-to-IGMP and PIM-to-MLD Message Translation Overview on page 239
- Configuring PIM-to-IGMP Message Translation on page 241
- Configuring PIM-to-MLD Message Translation on page 242

PIM-to-IGMP and PIM-to-MLD Message Translation Overview

This topic describes how the router translates Protocol Independent Multicast (PIM) join or prune messages into corresponding Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) report or leave messages. It also explains how you can use this feature to forward multicast traffic across PIM domains in certain network topologies.

In some network configurations, customers are unable to run PIM between the customer edge–facing PIM domain and the core-facing PIM domain, although PIM is running in sparse mode within each of these domains. Because PIM is not running between the domains, customers with this configuration cannot use PIM to forward multicast traffic across the domains. Instead, they might want to use IGMP to forward IPv4 multicast traffic, or MLD to forward IPv6 multicast traffic across the domains.

To enable the use of IGMP or MLD to forward multicast traffic across the PIM domains in such topologies, you can configure the rendezvous point (RP) router that resides between the edge domain and core domain to translate PIM join or prune messages received from PIM neighbors on downstream interfaces into corresponding IGMP or MLD report or leave messages. The router then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP router. As a result, this feature is sometimes referred to as *PIM-to-IGMP proxy* or *PIM-to-MLD proxy*.

To configure the RP router to translate PIM join or prune messages into IGMP report or leave messages, include the pim-to-igmp-proxy statement at the [edit routing-options multicast] hierarchy level. Similarly, to configure the RP router to translate PIM join or prune messages into MLD report or leave messages, include the pim-to-mld-proxy statement at the [edit routing-options multicast] hierarchy level. As part of the configuration, you must specify the full name of at least one, but not more than two,

upstream interfaces on which to enable the PIM-to-IGMP proxy or PIM-to-MLD proxy feature.

The following guidelines apply when you configure PIM-to-IGMP or PIM-to-MLD message translation:

- Make sure the router connecting the PIM edge domain and the PIM core domain is the static or elected RP router.
- Make sure the RP router is using the PIM sparse mode (PIM-SM) multicast routing protocol.
- When you configure an upstream interface, use the full logical interface specification (for example, ge-0/0/1.0) and not just the physical interface specification (ge-0/0/1).
- When you configure two upstream interfaces, the RP router transmits the same IGMP or MLD report messages and multicast traffic on both upstream interfaces. As a result, make sure that reverse-path forwarding (RPF) is running in the PIM-SM core domain to verify that multicast packets are received on the correct incoming interface and to avoid sending duplicate packets.
- The router transmits IGMP or MLD report messages on one or both upstream interfaces only for the first PIM join message that it receives among all of the downstream interfaces. Similarly, the router transmits IGMP or MLD leave messages on one or both upstream interfaces only if it receives a PIM prune message for the last downstream interface.
- Upstream interfaces support both local sources and remote sources.
- Multicast traffic received from an upstream interface is accepted as if it came from a host.

Related Topics

- Configuring PIM-to-IGMP Message Translation on page 241
- Configuring PIM-to-MLD Message Translation on page 242
- PIM Sparse Mode on page 342
- Example: Configuring PIM Sparse Mode on page 422

Configuring PIM-to-IGMP Message Translation

You can configure the rendezvous point (RP) routing device to translate PIM join or prune messages into corresponding IGMP report or leave messages. To do so, include the pim-to-igmp-proxy statement at the [edit routing-options multicast] hierarchy level:

```
[edit routing-options multicast]
pim-to-igmp-proxy {
    upstream-interface [ interface-names ];
}
```

Enabling the routing device to perform PIM-to-IGMP message translation, also referred to as PIM-to-IGMP proxy, is useful when you want to use IGMP to forward IPv4 multicast traffic between a PIM sparse mode edge domain and a PIM sparse mode core domain in certain network topologies.

Before you begin configuring PIM-to-IGMP message translation:

- Make sure the routing device connecting the PIM edge domain and the PIM core domain is the static or elected RP routing device.
- Make sure the PIM sparse mode (PIM-SM) routing protocol is running on the RP routing device.
- If you plan to configure two upstream interfaces, make sure that reverse-path forwarding (RPF) is running in the PIM-SM core domain. Because the RP router transmits the same IGMP messages and multicast traffic on both upstream interfaces, you need to run RPF to verify that multicast packets are received on the correct incoming interface and to avoid sending duplicate packets.

To configure the RP routing device to translate PIM join or prune messages into corresponding IGMP report or leave messages:

1. Include the pim-to-igmp-proxy statement, specifying the names of one or two logical interfaces to function as the upstream interfaces on which the routing device transmits IGMP report or leave messages.

   The following example configures PIM-to-IGMP message translation on a single upstream interface, ge-0/1/0.1.

   ```
   [edit routing-options multicast]
   user@host# set pim-to-igmp-proxy upstream-interface ge-0/1/0.1
   ```

   The following example configures PIM-to-IGMP message translation on two upstream interfaces, ge-0/1/0.1 and ge-0/1/0.2. You must include the logical interface names within square brackets ([ ]) when you configure a set of two upstream interfaces.

   ```
   [edit routing-options multicast]
   user@host# set pim-to-igmp-proxy upstream-interface [ge-0/1/0.1 ge-0/1/0.2]
   ```

2. Use the show multicast pim-to-igmp-proxy command to display the PIM-to-IGMP proxy state (enabled or disabled) and the name or names of the configured upstream interfaces.

```
user@host# run show multicast pim-to-igmp-proxy

Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

Related Topics

- PIM-to-IGMP and PIM-to-MLD Message Translation Overview on page 239
- pim-to-igmp-proxy
- upstream-interface

Configuring PIM-to-MLD Message Translation

You can configure the rendezvous point (RP) routing device to translate PIM join or prune messages into corresponding MLD report or leave messages. To do so, include the pim-to-mld-proxy statement at the [edit routing-options multicast] hierarchy level:

```
[edit routing-options multicast]
pim-to-mld-proxy {
   upstream-interface [ interface-names ];
}
```

Enabling the routing device to perform PIM-to-MLD message translation, also referred to as PIM-to-MLD proxy, is useful when you want to use MLD to forward IPv6 multicast traffic between a PIM sparse mode edge domain and a PIM sparse mode core domain in certain network topologies.

Before you begin configuring PIM-to-MLD message translation:

- Make sure the routing device connecting the PIM edge domain and the PIM core domain is the static or elected RP routing device.
- Make sure the PIM sparse mode (PIM-SM) routing protocol is running on the RP routing device.
- If you plan to configure two upstream interfaces, make sure that reverse-path forwarding (RPF) is running in the PIM-SM core domain. Because the RP routing device transmits the same MLD messages and multicast traffic on both upstream interfaces, you need to run RPF to verify that multicast packets are received on the correct incoming interface and to avoid sending duplicate packets.

To configure the RP routing device to translate PIM join or prune messages into corresponding MLD report or leave messages:

1. Include the pim-to-mld-proxy statement, specifying the names of one or two logical interfaces to function as the upstream interfaces on which the router transmits MLD report or leave messages.

   The following example configures PIM-to-MLD message translation on a single upstream interface, ge-0/5/0.1.

   ```
   [edit routing-options multicast]
   user@host# set pim-to-mld-proxy upstream-interface ge-0/5/0.1
   ```

The following example configures PIM-to-MLD message translation on two upstream interfaces, ge-0/5/0.1 and ge-0/5/0.2. You must include the logical interface names within square brackets ([ ]) when you configure a set of two upstream interfaces.

```
[edit routing-options multicast]
user@host# set pim-to-mld-proxy upstream-interface [ge-0/5/0.1 ge-0/5/0.2]
```

2. Use the show multicast pim-to-mld-proxy command to display the PIM-to-MLD proxy state (enabled or disabled) and the name or names of the configured upstream interfaces.

```
user@host# run show multicast pim-to-mld-proxy

Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

Related Topics

- PIM-to-IGMP and PIM-to-MLD Message Translation Overview on page 239
- pim-to-mld-proxy
- upstream-interface

Chapter 28

Summary of Multicast Routing Options Configuration Statements

The following sections explain each of the multicast routing options configuration statements. The statements are organized alphabetically.

asm-override-ssm

Syntax asm-override-ssm;

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

Release Information Statement introduced in JUNOS Release 9.4.
Statement introduced in JUNOS Release 9.5 for EX Series switches.

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 Topics
- Configuring Source-Specific Multicast Groups on page 204

backup-pe-group

Syntax

```
backup-pe-group group-name {
    backups [ addresses ];
    local-address address;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in JUNOS Release 9.0.
Statement introduced in JUNOS Release 9.5 for EX Series switches.

Description

Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options

group-name—Name of the group for PE backups.

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 Topics

- Configuring Ingress PE Redundancy on page 235

backups

Syntax backups [*addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit logical-systems *logical-system-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-options multicast backup-pe-group *group-name*]

Release Information Statement introduced in JUNOS Release 9.0.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the address of backup PEs for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options *addresses*—Addresses of other PEs in the backup group.

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

Related Topics
- Configuring Ingress PE Redundancy on page 235

bandwidth

Syntax bandwidth (*bps* | adaptive);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map],
[edit logical-systems *logical-system-name* routing-options multicast flow-map],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map],
[edit routing-options multicast flow-map]

Release Information Statement introduced in JUNOS Release 8.3.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the bandwidth property for multicast flow maps.

Options adaptive—Specify that the bandwidth is measured for the flows that are matched by the flow map.

pbs—Bandwidth, in bits per second, for the flow map.
Range: 0 through any amount of bandwidth
Default: 2 Mbps

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

Related Topics
- Defining Bandwidth for Multicast Flows on page 218

flow-map

Syntax

```
flow-map flow-map-name {
   bandwidth (bps | adaptive);
   forwarding-cache {
      timeout (never | minutes);
   }
   policy [ policy-names ];
   redundant-sources [ addresses ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in JUNOS Release 8.2.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure multicast flow maps.

Options

flow-map-name—Name of the flow-map.

The remaining statements are explained separately.

Required Privilege Level

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

Related Topics

- Creating a Multicast Flow Map on page 211

forwarding-cache (Flow Maps)

Syntax

```
forwarding-cache {
   timeout (minutes | never );
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information

Statement introduced in JUNOS Release 8.2.

Description

Configure multicast forwarding cache properties for the flow map.

Required Privilege Level

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

forwarding-cache (Multicast)

Syntax

```
forwarding-cache {
   threshold suppress value <reuse value>;
   timeout minutes;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure multicast forwarding cache properties. These properties include threshold suppression and reuse limits and timeout values.

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 Topics

- Configuring General Multicast Forwarding Cache Properties on page 229

interface (Routing Options)

Syntax

```
interface interface-names {
   maximum-bandwidth bps;
   no-qos-adjust;
   reverse-oif-mapping {
      no-qos-adjust;
   }
   subscriber-leave-timer seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in JUNOS Release 8.3.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Define the maximum bandwidth for an interface on which you want to apply bandwidth management.

Options

interface-name—Names of the physical or logical interface. For details about specifying interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

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.

interface (Scoping)

Syntax interface [*interface-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit logical-systems *logical-system-name* routing-options multicast scope *scope-name*],
[edit routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit routing-options multicast scope *scope-name*]

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

Description Configure the set of interfaces for multicast scoping.

Options *interface-names*—Names of the interfaces to scope. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify all. For details about specifying interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

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

Related Topics
- Creating a Named Scope for Multicast Scoping on page 194

local-address

Syntax local-address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit logical-systems *logical-system-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-options multicast backup-pe-group *group-name*]

Release Information Statement introduced in JUNOS Release 9.0.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the address of the local PE for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

Options *address*—Address of local PEs in the backup group.

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

Related Topics
- Configuring Ingress PE Redundancy on page 235

maximum-bandwidth

Syntax maximum-bandwidth *bps*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name*]

Release Information Statement introduced in JUNOS Release 8.3.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the multicast bandwidth for the interface.

Options *bps*—Bandwidth rate, in bits per second, for the multicast interface.
Range: 0 through any amount of bandwidth

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

Related Topics
- Defining Interface Bandwidth Maximums on page 218

multicast

Syntax

```
multicast {
   asm-override-ssm;
   backup-pe-group group-name {
      backups [ addresses ];
      local-address address;
   }
   flow-map flow-map-name {
      bandwidth (bps | adaptive);
      forwarding-cache {
         timeout (never | minutes);
      }
      policy [ policy-names ];
      redundant-sources [ addresses ];
   }
   forwarding-cache {
      threshold suppress value <reuse value>;
      timeout minutes;
   }
   interface interface-name {
      maximum-bandwidth bps;
      no-qos-adjust;
      reverse-oif-mapping {
         no-qos-adjust;
      }
      subscriber-leave-timer seconds;
   }
   pim-to-igmp-proxy {
      upstream-interface [ interface-names ];
   }
   pim-to-mld-proxy {
      upstream-interface [ interface-names ];
   }
   rpf-check-policy [ policy-names ];
   scope scope-name {
      interface [ interface-names ];
      prefix destination-prefix;
   }
   scope-policy [ policy-names ];
   ssm-groups [ addresses ];
   ssm-map ssm-map-name {
      policy [ policy-names ];
      source [ addresses ];
   }
   traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <disable>;
   }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* routing-options],

[edit dynamic-profiles *profile-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

NOTE: You cannot apply a scope policy to a specific routing instance. That is, all scoping policies are applied to all routing instances. However, the scope statement does apply individually to a specific routing instance.

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.
interface and maximum-bandwidth statements introduced in JUNOS Release 8.3.
interface and maximum-bandwidth statements introduced in JUNOS Release 9.0 for EX Series switches.
Statement added to [edit dynamic-profiles routing-options] and [edit dynamic-profiles *profile-name* routing-instances *routing-instance-name* routing-options] hierarchy levels in JUNOS Release 9.6.

Description

Configure multicast routing options properties.

The remaining statements are explained separately.

Required Privilege Level

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

Related Topics

- Configuring General Multicast Forwarding Cache Properties on page 229
- Configuring Multicast Forwarding Cache Properties for Flow Maps on page 230
- Configuring Source-Specific Multicast Groups on page 204
- *JUNOS Multicast Protocols Configuration Guide*

no-qos-adjust

Syntax

no-qos-adjust;

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-options multicast interface *interface-name*]
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name* reverse-oif-mapping],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name* reverse-oif-mapping],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name* reverse-oif-mapping],
[edit routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name* reverse-oif-mapping]

Release Information

Statement introduced in JUNOS Release 9.5.
Statement introduced in JUNOS Release 9.5 for EX Series switches.
Statement added to [edit dynamic-profiles routing-options multicast interface *interface-name*], [edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*], [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*], and [edit routing-options multicast interface *interface-name*] hierarchy levels in JUNOS Release 9.6.

Description

Disable hierarchical bandwidth adjustment for all subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.

Required Privilege Level

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

Related Topics

- Managing Subscriber Overcommitment on page 226

pim-to-igmp-proxy

Syntax

```
pim-to-igmp-proxy {
   upstream-interface [ interface-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in JUNOS Release 9.6.
Statement introduced in JUNOS Release 9.6 for EX Series switches.

Description

Configure the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Internet Group Management Protocol (IGMP) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the pim-to-igmp-proxy statement enables you to use IGMP to forward IPv4 multicast traffic across the PIM sparse mode domains.

The remaining statement is explained separately.

Required Privilege Level

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

Related Topics

- Configuring PIM-to-IGMP Message Translation on page 241

pim-to-mld-proxy

Syntax

```
pim-to-mld-proxy {
   upstream-interface [ interface-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in JUNOS Release 9.6.
Statement introduced in JUNOS Release 9.6 for EX Series switches.

Description

Configure the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Multicast Listener Discovery (MLD) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the pim-to-mld-proxy statement enables you to use MLD to forward IPv6 multicast traffic across the PIM sparse mode domains.

The remaining statement is explained separately.

Required Privilege Level

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

Related Topics

- Configuring PIM-to-MLD Message Translation on page 242

policy (Flow Maps)

Syntax policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information Statement introduced in JUNOS Release 8.2.

Description Configure a flow map policy.

Options *policy-names*—Name of one or more policies for flow mapping.

Required Privilege Level routing—To view this statement in the configuration.

policy (SSM Maps)

Syntax policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-options multicast ssm-map *ssm-map-name*]

Release Information Statement introduced in JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Apply one or more policies to an SSM map.

Options *policy-names*—Name of one or more policies for SSM mapping.

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

Related Topics
- Example: Configuring SSM Mapping on page 208

prefix

Syntax prefix *destination-prefix*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit logical-systems *logical-system-name* routing-options multicast scope *scope-name*],
[edit routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit routing-options multicast scope *scope-name*]

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

Description Configure the prefix for multicast scopes.

Options *destination-prefix*—Address range for the multicast scope.

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

Related Topics
- Creating a Named Scope for Multicast Scoping on page 194

redundant-sources

Syntax redundant-sources [*addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information Statement introduced in JUNOS Release 8.3.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure a list of redundant sources for multicast flows defined by a flow map.

Options *addresses*—List of IPv4 or IPv6 addresses for use as redundant (backup) sources for multicast flows defined by a flow map.

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

Related Topics
- Specifying Redundant Flow Sources on page 212

reverse-oif-mapping

Syntax

```
reverse-oif-mapping {
   no-qos-adjust;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name*]

Release Information

Statement introduced in JUNOS Release 9.2.
Statement introduced in JUNOS Release 9.2 for EX Series switches.
The no-qos-adjust statement added in JUNOS Release 9.5.
The no-qos-adjust statement introduced in JUNOS Release 9.5 for EX Series switches.

Description

Enable the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD request it receives over the multicast VLAN.

The remaining statement is explained separately.

Required Privilege Level

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

Related Topics

- Managing Subscriber Overcommitment on page 226

rpf-check-policy

Syntax

rpf-check-policy [*policy-names*];

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in JUNOS Release 8.1.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.

Options

policy-names—Name of one or more multicast RPF check policies.

Required Privilege Level

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

Related Topics

- Configuring RPF Policies on page 200

scope

Syntax

```
scope scope-name {
   interface [ interface-names ];
   prefix destination-prefix;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure multicast scoping.

Options

scope-name—Name of the multicast scope.

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 Topics

- Creating a Named Scope for Multicast Scoping on page 194

scope-policy

Syntax scope-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-options multicast]

NOTE: You can configure a scope policy at these two hierarchy levels only. You cannot apply a scope policy to a specific routing instance, because all scoping policies are applied to all routing instances. However, you can apply the scope statement to a specific routing instance at the [edit routing-instances *routing-instance-name* routing-options multicast] or [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast] hierarchy level.

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

Description Apply policies for scoping. The policy must be correctly configured at the edit policy-options policy-statement hierarchy level.

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

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

Related Topics

- scope
- Using a Scope Policy for Multicast Scoping on page 195

source

Syntax source [*addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-options multicast ssm-map *ssm-map-name*]

Release Information Statement introduced in JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Specify IPv4 or IPv6 source addresses for an SSM map.

Options *addresses*—IPv4 or IPv6 source addresses.

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

Related Topics
- Example: Configuring SSM Mapping on page 208

ssm-groups

Syntax ssm-groups [*ip-addresses*];

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

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

Description Configure additional source-specific multicast (SSM) groups.

Options *ip-addresses*—List of one or more additional SSM group addresses separated by a space.

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

Related Topics
- Configuring Source-Specific Multicast Groups on page 204

ssm-map

Syntax

```
ssm-map ssm-map-name {
    policy [ policy-names ];
    source [ addresses ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure SSM mapping.

Options

ssm-map-name—Name of the SSM map.

The remaining statements are explained separately.

Required Privilege Level

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

Related Topics

- Example: Configuring SSM Mapping on page 208

subscriber-leave-timer

Syntax subscriber-leave-timer *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name*]

Release Information Statement introduced in JUNOS Release 9.2.
Statement introduced in JUNOS Release 9.2 for EX Series switches.

Description Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message.

Options *seconds*—Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message. Specifying a value of 0 results in an immediate update; this is the same as if the statement were not configured.
Range: 0 through 30
Default: 0 seconds

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

Related Topics
- Managing Subscriber Overcommitment on page 226

threshold

Syntax threshold suppress *value* <reuse *value*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface],
[edit logical-systems *logical-system-name* routing-options multicast interface],
[edit routing-instances *routing-instance-name* routing-options multicast interface],
[edit routing-options multicast interface]

Release Information Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.2 for EX Series switches.

Description Configure the suppression and reuse thresholds for multicast forwarding cache limits.

Options reuse *value*—(Optional) Value to begin creating new multicast forwarding cache entries. This value is optional. If configured, this number must be less than the suppress value.
Range: 1 through 200,000

suppress *value*—Value to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number must be greater than the reuse value.
Range: 1 through 200,000

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

Related Topics
- Examples: Configuring Multicast Forwarding Cache Properties on page 231

timeout (Flow Maps)

Syntax timeout (|);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information Statement introduced in JUNOS Release 8.2.

Description Configure the timeout value for multicast forwarding cache entries associated with the flow map.

Options *minutes*—Length of time that the forwarding cache entry remains active.
Range: 1 through 720

never—Specifies that the forwarding cache entry always remains active.

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

timeout (Multicast)

Syntax timeout *minutes*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast forwarding-cache],
[edit logical-systems *logical-system-name* routing-options multicast forwarding-cache],
[edit routing-instances *routing-instance-name* routing-options multicast forwarding-cache],
[edit routing-options multicast forwarding-cache]

Release Information Statement introduced in JUNOS Release 8.2.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the timeout value for multicast forwarding cache entries.

Options *minutes*—Length of time that the forwarding cache limit remains active.
Range: 1 through 720

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

Related Topics
- Configuring General Multicast Forwarding Cache Properties on page 229

upstream-interface

Syntax

upstream-interface [*interface-names*];

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast pim-to-igmp-proxy],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast pim-to-mld-proxy],
[edit logical-systems *logical-system-name* routing-options multicast pim-to-igmp-proxy],
[edit logical-systems *logical-system-name* routing-options multicast pim-to-mld-proxy],
[edit routing-instances *routing-instance-name* routing-options multicast pim-to-igmp-proxy],
[edit routing-instances *routing-instance-name* routing-options multicast pim-to-mld-proxy],
[edit routing-options multicast pim-to-igmp-proxy],
[edit routing-options multicast pim-to-mld-proxy]

Release Information

Statement introduced in JUNOS Release 9.6.
Statement introduced in JUNOS Release 9.6 for EX Series switches.

Description

Configure at least one, but not more than two, upstream interfaces on the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain. The RP routing device translates PIM join or prune messages into corresponding IGMP report or leave messages (if you include the pim-to-igmp-proxy statement), or into corresponding MLD report or leave messages (if you include the pim-to-mld-proxy statement). The routing device then proxies the IGMP or MLD report or leave messages to one or both upstream interfaces to forward IPv4 multicast traffic (for IGMP) or IPv6 multicast traffic (for MLD) across the PIM domains.

Options

interface-names—Names of one or two upstream interfaces to which the RP routing device proxies IGMP or MLD report or leave messages for transmission of multicast traffic across PIM domains. You can specify a maximum of two upstream interfaces on the RP routing device. To configure a set of two upstream interfaces, specify the full interface names, including all physical and logical address components, within square brackets ([ ]). For details about specifying interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

Required Privilege Level

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

Related Topics

- Configuring PIM-to-IGMP Message Translation on page 241
- Configuring PIM-to-MLD Message Translation on page 242

Part 9
Multicast Snooping

- Multicast Snooping Overview on page 275
- Multicast Snooping Configuration Guidelines on page 277
- Summary of Multicast Snooping Configuration Statements on page 283
- IGMP Snooping Overview on page 289
- IGMP Snooping Configuration Guidelines on page 293
- Summary of IGMP Snooping Configuration Statements on page 299

Chapter 29

Multicast Snooping Overview

Generally, multicast snooping is a way for a Layer 2 device to "snoop" at the Layer 3 packet content to determine which actions are to be taken to process or forward a frame. There are specific forms of snooping, such as IGMP snooping or PIM snooping. In all cases, snooping involves a device configured to function at Layer 2 having access to Layer 3 (packet) information. Snooping makes multicasting more efficient in these devices.

For information about Layer 2 and multicast addressing, see "Layer 2 Frames and IPv4 Multicast Addresses" on page 21. For more information about multicast snooping, see "Multicast Snooping Overview" on page 23.

Chapter 30

Multicast Snooping Configuration Guidelines

This chapter describes the following tasks for configuring multicast snooping:

- Configuring Multicast Snooping on page 277
- Configuring Forwarding Cache Snooping Options on page 278
- Configuring Flood Groups for Snooping on page 279
- Configuring Graceful Restart for Snooping on page 279
- Ignoring Spanning Tree Topology Change Messages on page 280

Configuring Multicast Snooping

To configure the general multicast snooping parameters for MX Series routers, include the multicast-snooping-options statement:

```
multicast-snooping-options {
   flood-groups [ ip-addresses ];
   forwarding-cache {
      threshold suppress value <reuse value>;
   }
   graceful-restart <restart-duration seconds>;
   ignore-stp-topology-change;
   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 bridge-domains *bridge-domain-name*]
- [edit routing-instances *routing-instance-name*]
- [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*]

By default, multicast snooping is disabled. You can enable multicast snooping in VPLS or virtual switch instance types in the instance hierarchy or in one or more bridge domains.

If there are multiple bridge domains configured under a VPLS or virtual switch instance, the multicast snooping options configured at the instance level apply to all the bridge domains. Multicast snooping options configured at the bridge domain level only apply to that particular bridge domain. The options configured at the bridge domain take precedence over the options configured at the instance level.

NOTE: The ignore-stp-topology-change statement is supported for the virtual-switch routing instance type only and is not supported under the [edit logical-systems] hierarchy.

Related Topics

- Configuring IGMP Snooping on page 293
- Configuring IGMP Snooping Proxy Mode on page 294
- Configuring IGMP Snooping Immediate Leave on page 294
- Configuring General IGMP Snooping Parameters on page 295
- Configuring IGMP Snooping Interfaces on page 297
- Configuring VLAN-Specific IGMP Snooping Parameters on page 298
- Tracing IGMP Snooping Operations on page 298

Configuring Forwarding Cache Snooping Options

You can configure threshold values on the forwarding cache to suppress (suspend) snooping when the cache entries reach a certain maximum and reuse the cache when the number falls to another threshold value. By default, no threshold values are enabled on the router.

You can specify a value for the threshold to suppress new multicast forwarding cache entries and an optional reuse value for the threshold at which the router begins to create new multicast forwarding cache entries. The range for both is from 1 through 200,000. If configured, the reuse value must be less than the suppression threshold value. The suppression value is mandatory. If you do not specify the optional reuse value, then the number of multicast forwarding cache entries is limited to the suppression value. A new entry is created as soon as the number of multicast forwarding cache entries falls below the suppression value.

To configure threshold values for a forwarding cache, assign values to either the suppress or reuse threshold and include the forwarding-cache statement:

```
forwarding-cache {
   threshold suppress value <reuse value>;
}
```

You can include these statements at the following hierarchy levels:

- [edit bridge-domains *bridge-domain-name* multicast-snooping-options]
- [edit routing-instances *routing-instance-name* multicast-snooping-options]
- [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options]

Configuring Flood Groups for Snooping

Flood groups are lists of multicast addresses. By default, the router does not track flood group addresses.

To add IP addresses to the flood groups, include the flood-groups statement:

```
flood-groups [ ip-addresses ];
```

You can include this statement at the following hierarchy levels:

- [edit bridge-domains *bridge-domain-name* multicast-snooping-options]
- [edit routing-instances *routing-instance-name* multicast-snooping-options]
- [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options]

Configuring Graceful Restart for Snooping

You can configure graceful restart and duration time for multicast snooping. If graceful restart for multicast snooping is disabled, snooping information is lost after a Routing Engine restart.

By default, the graceful restart duration is 180 seconds (3 minutes). You can set this value between 0 and 300 seconds. If you set the duration to 0, graceful restart is effectively disabled. Set this value slightly larger than the IGMP query response interval. For more information about the IGMP query response interval, see query-response-interval.

To configure graceful restart for multicast snooping, include the graceful-restart statement:

```
graceful-restart <restart-duration seconds>;
```

You can include this statement at the following hierarchy levels:

- [edit bridge-domains *bridge-domain-name* multicast-snooping-options]
- [edit routing-instances *routing-instance-name* multicast-snooping-options]
- [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options]

Ignoring Spanning Tree Topology Change Messages

By default the IGMP snooping process on an MX Series router detects interface state changes made by any of the spanning tree protocols (STPs).

In a VPLS multihoming environment where two PE routers are connected to two interconnected CE routers and STP root protection is enabled on the PE routers, one of the PE router interfaces is in forwarding state and the other is in blocking state.

If the link interconnecting the two CE routers fails, the PE router interface in blocking state transitions to the forwarding state.

The PE router interface does not wait to receive membership reports in response to the next general or group-specific query. Instead, the IGMP snooping process sends a general query message toward the CE router. The hosts connected to the CE router reply with reports for all groups they are interested in.

When the link interconnecting the two CE routers is restored, the original spanning tree state on both PR routers is restored. The forwarding PE receives a spanning tree topology change message and sends a general query message towards the CE router to immediately reconstruct the group membership state.

The default behavior for the IGMP snooping process on an MX Series router can be changed to ignore messages about the spanning tree topology state change.

To ignore the spanning tree interface state, include the ignore-stp-topology-change statement:

```
ignore-stp-topology-change;
```

You can include this statement at the following hierarchy levels:

- [edit bridge-domains *bridge-domain-name* multicast-snooping-options]
- [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options]

NOTE: The ignore-stp-topology-change statement is supported for the virtual-switch routing instance type only.

Chapter 31

Summary of Multicast Snooping Configuration Statements

The following sections explain each of the general multicast snooping configuration statements. The statements are organized alphabetically.

flood-groups

Syntax flood-groups [*ip-addresses*];

Hierarchy Level [edit bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* multicast-snooping-options]

Release Information Statement introduced in JUNOS Release 8.5.

Description Establish a list of flood group addresses for multicast snooping.

Options *ip-addresses*—List of IP addresses subject to flooding.

Usage Guidelines See "Configuring Flood Groups for Snooping" on page 279.

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

forwarding-cache

Syntax

```
forwarding-cache {
   threshold suppress value <reuse value>;
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* multicast-snooping-options]

Release Information

Statement introduced in JUNOS Release 8.5.

Description

Establish multicast snooping forwarding cache parameter values.

Options

The statements are explained separately.

Usage Guidelines

See “Configuring Forwarding Cache Snooping Options” on page 278.

Required Privilege Level

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

graceful-restart

Syntax graceful-restart <restart-duration *seconds*>;

Hierarchy Level [edit bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* multicast-snooping-options]

Release Information Statement introduced in JUNOS Release 9.2.

Description Establish the graceful restart duration for multicast snooping. You can set this value between 0 and 300 seconds. If you set the duration to 0, graceful restart is effectively disabled. Set this value slightly larger than the IGMP query response interval.

Default 180 seconds

Usage Guidelines See "Configuring Graceful Restart for Snooping" on page 279.

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

Related Topics query-response-interval

ignore-stp-topology-change

Syntax ignore-stp-topology-change;

Hierarchy Level [edit bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options]

Release Information Statement introduced in JUNOS Release 9.5.

Description Ignore messages about spanning tree topology changes. This statement is supported for the virtual-switch routing instance type only.

Usage Guidelines See "Ignoring Spanning Tree Topology Change Messages" on page 280.

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

multicast-snooping-options

Syntax

```
multicast-snooping-options {
   flood-groups [ ip-addresses ];
   forwarding-cache {
      threshold suppress value <reuse value>;
   }
   graceful-restart <restart-duration seconds>;
   ignore-stp-topology-change;
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*],
[edit routing-instances *routing-instance-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*]

Release Information

Statement introduced in JUNOS Release 8.5.

Description

Establish multicast snooping option values.

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 Topics

- Configuring Multicast Snooping on page 277
- Configuring Forwarding Cache Snooping Options on page 278
- Configuring Flood Groups for Snooping on page 279
- Configuring Graceful Restart for Snooping on page 279
- Ignoring Spanning Tree Topology Change Messages on page 280

threshold

Syntax threshold suppress *value* <reuse *value*>;

Hierarchy Level [edit bridge-domains *bridge-domain-name* multicast-snooping-options forwarding-cache],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options forwarding-cache],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options forwarding-cache],
[edit routing-instances *routing-instance-name* multicast-snooping-options forwarding-cache],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options forwarding-cache]

Release Information Statement introduced in JUNOS Release 8.5.

Description Configure the suppression and reuse thresholds for multicast snooping forwarding cache limits.

Options suppress *value*—Value to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number must be greater than the reuse value.
Range: 1 through 200,000

reuse *value*—(Optional) Value to begin creating new multicast forwarding cache entries. If configured, this number must be less than the suppress value.
Range: 1 through 200,000

Usage Guidelines See "Configuring Forwarding Cache Snooping Options" on page 278.

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

Chapter 32

IGMP Snooping Overview

This chapter presents the following topics about IGMP snooping on MX Series routers:

- Introduction to IGMP Snooping on page 289
- IGMP Snooping Interfaces and Forwarding on page 290
- IGMP Snooping and Proxies on page 291
- Multicast-Router Interfaces and IGMP Snooping Proxy Mode on page 292
- Host-Side Interfaces and IGMP Snooping Proxy Mode on page 292
- IGMP Snooping and Bridge Domains on page 292

Introduction to IGMP Snooping

Snooping is a general way for a Layer 2 device, such as Juniper Networks MX Series Ethernet Services Routers, to implement a series of procedures to “snoop” at the Layer 3 packet content to determine which actions are to be taken to process or forward a frame. More specific forms of snooping, such as Internet Group Membership Protocol (IGMP) snooping or Protocol Independent Multicast (PIM) snooping, are used with multicast.

Layer 2 devices (LAN switches or bridges) handle multicast packets and the frames that contain them much in the same way the Layer 3 devices (routers) handle broadcasts. So a Layer 2 switch processes an arriving frame having a multicast destination media access control (MAC) address by forwarding a copy of the packet (frame) onto each of the other network interfaces of the switch that are in a forwarding state.

However, this approach (sending multicast frames everywhere the device can) is not the most efficient use of network bandwidth, particularly for IPTV applications. IGMP snooping functions by “snooping” at the IGMP packets received by the switch interfaces and building a multicast database similar to that a multicast router builds in a Layer 3 network. Using this database, the switch can forward multicast traffic only onto downstream interfaces with interested receivers, and this technique allows more efficient use of network bandwidth.

You configure IGMP snooping for each bridge on the router. A bridge instance without qualified learning has just one learning domain. For a bridge instance with qualified learning, snooping will function separately within each learning domain in the bridge. That is, IGMP snooping and multicast forwarding will proceed independently in each learning domain in the bridge.

This chapter discusses only bridge instances without qualified learning (those forming one learning domain on the device). Therefore, all the interfaces mentioned in this chapter are logical interfaces of the bridge or VPLS instance.

Several related concepts are important when discussing IGMP snooping:

- Bridge or VPLS instance interfaces are either multicast-router interfaces or host-side interfaces.
- IGMP snooping supports proxy mode or without-proxy mode.

NOTE: When integrated routing and bridging (IRB) is used, if the router is an IGMP querier, any leave message received on any Layer 2 interface will cause a group-specific query on all Layer 2 interfaces (as a result of this practice, some corresponding reports might be received on all Layer 2 interfaces). However, if some of the Layer 2 interfaces are also router (Layer 3) interfaces, reports and leaves from other Layer 2 interfaces will not be forwarded on those interfaces.

If an IRB interface is used as an outgoing interface in a multicast forwarding cache entry (as determined by the routing process), then the output interface list will be expanded into a subset of the Layer 2 interface in the corresponding bridge. The subset is based on the snooped multicast membership information, according to the multicast forwarding cache entry installed by the snooping process for the bridge.

If no snooping is configured, the IRB output interface list is expanded to all Layer 2 interfaces in the bridge.

For more information about Layer 2 and multicast addressing, see “Layer 2 Frames and IPv4 Multicast Addresses” on page 21. For more information about multicast snooping in general, see “Multicast Snooping Overview” on page 23.

IGMP Snooping Interfaces and Forwarding

IGMP snooping divides the device interfaces into multicast-router interfaces and host-side interfaces. A multicast-router interface is an interface in the direction of a multicasting router. An interface on the bridge is considered a multicast-router interface if it meets at least one of the following criteria:

- It is statically configured as a multicast-router interface in the bridge instance.
- IGMP queries are being received on the interface.

All other interfaces that are not multicast-router interfaces are considered host-side interfaces.

Any multicast traffic received on a bridge interface with IGMP snooping configured will be forwarded according to following rules:

- Any IGMP packet is sent to the Routing Engine for snooping processing.
- Other multicast traffic with destination address 224.0.0/24 is flooded onto all other interfaces of the bridge.

- Other multicast traffic is sent to all the multicast-router interfaces but only to those host-side interfaces that have hosts interested in receiving that multicast group.

IGMP Snooping and Proxies

Without a proxy arrangement, IGMP snooping does not generate or introduce queries and reports. It will only "snoop" reports received from all of its interfaces (including multicast-router interfaces) to build its state and group (S,G) database.

Without a proxy, IGMP messages are processed as follows:

- Query—All general and group-specific IGMP query messages received on a multicast-router interface will be forwarded to all other interfaces (both multicast-router interfaces and host-side interfaces) on the bridge.
- Report—IGMP reports received on any interface of the bridge are forwarded towards other multicast-router interfaces. The receiving interface is added as an interface for that group if a multicast routing entry exists for this group. Also, a group timer is set for the group on that interface. If this timer expires (that is, there was no report for this group during the IGMP group timer period), then the interface is removed as an interface for that group.
- Leave—Any IGMP leave message received on any interface of the bridge.

Proxy snooping reduces the number of IGMP reports sent towards an IGMP router.

NOTE: With proxy snooping configured, an IGMP router is not able to perform host tracking.

As proxy for its host-side interfaces, IGMP snooping in proxy mode replies to the queries it receives from an IGMP router on a multicast-router interface. On the host-side interfaces, IGMP snooping in proxy mode behaves as an IGMP router and sends general and group-specific queries on those interfaces.

NOTE: Only group-specific queries are generated by IGMP snooping directly; general queries received from the multicast-router interfaces are flooded to host-side interfaces.

All the queries generated by IGMP snooping are sent using 0.0.0.0 as source address. Also, all reports generated by IGMP snooping are sent with 0.0.0.0 as source address unless there is a configured source address to use.

Proxy mode functions differently on multicast-router interfaces than it does on host-side interfaces.

Multicast-Router Interfaces and IGMP Snooping Proxy Mode

On multicast-router interfaces, in response to IGMP queries, IGMP snooping in proxy mode sends reports containing aggregate information on groups learned on all host-side interfaces of the bridge.

Besides replying to queries, the IGMP snooping in proxy mode forwards all queries, reports, and leaves received on a multicast-router interface to other multicast-router interfaces. IGMP snooping keeps the membership information learned on this interface but does not send a group-specific query for leave messages received on this interface. It simply times out the groups learned on this interface if there are no reports for the same group within the timer duration.

NOTE: For the hosts on all the multicast-router interfaces, it is the IGMP router, not the IGMP snooping proxy, that generates general and group-specific queries.

Host-Side Interfaces and IGMP Snooping Proxy Mode

No reports are sent on host-side interfaces by IGMP snooping in proxy mode. IGMP snooping processes reports received on these interfaces and sends group-specific queries onto host-side interfaces when it receives a leave message on the interface. Host-side interfaces do not generate periodic general queries, but forwards or floods general queries received from multicast-router interfaces.

If a group is removed from a host-side interface and this was the last host-side interface for that group, a leave is sent to the multicast-router interfaces. If a group report is received on a host-side interface and this was the first host-side interface for that group, a report is sent to all multicast-router interfaces.

IGMP Snooping and Bridge Domains

IGMP snooping on a VLAN is only allowed for the legacy vlan-id all case. In other cases, there is a specific bridge domain configuration that determines the VLAN-specific configuration for IGMP snooping.

Chapter 33

IGMP Snooping Configuration Guidelines

This chapter describes the following tasks for configuring IGMP snooping:

- Configuring IGMP Snooping on page 293
- Configuring IGMP Snooping Proxy Mode on page 294
- Configuring IGMP Snooping Immediate Leave on page 294
- Configuring General IGMP Snooping Parameters on page 295
- Configuring IGMP Snooping Interfaces on page 297
- Configuring VLAN-Specific IGMP Snooping Parameters on page 298
- Tracing IGMP Snooping Operations on page 298

Configuring IGMP Snooping

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

```
igmp-snooping {
    immediate-leave;
    interface interface-name {
        group-limit limit;
        host-only-interface;
        immediate-leave;
        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;
    vlan vlan-id {
        immediate-leave;
        interface interface-name {
            group-limit limit;
```

```
            host-only-interface;
            immediate-leave;
            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]
- [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols]

By default, IGMP snooping is not enabled. Statements configured at the VLAN level apply only to that particular VLAN.

Configuring IGMP Snooping Proxy Mode

To configure IGMP snooping proxy mode, include the proxy statement:

```
proxy {
    source-address ip-address;
}
```

You can include this statement at the following hierarchy levels:

- [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 igmp-snooping vlan *vlan-id*]

Configuring IGMP Snooping Immediate Leave

You can configure a router running IGMP version 2 (IGMPv2) snooping so that, after the router receives a leave group membership message from a host associated with the interface, the router immediately removes the group membership from the

interface and suppresses the sending of any group-specific queries for the multicast group.

When IGMP snooping is enabled on a router running IGMP version 3 (IGMPv3) snooping, after the router receives a report with the type BLOCK_OLD_SOURCES, the router suppresses the sending of group-and-source queries but relies on the JUNOS Software's host-tracking mechanism to determine whether or not it removes a particular source group membership from the interface.

NOTE: When configuring this feature on IGMPv2 interfaces, ensure that the IGMP interface has only one IGMP host connected. If more than one IGMPv2 host is connected to a LAN through the same interface, and one host sends a leave message, the router removes all hosts on the interface from the multicast group. The router loses contact with the hosts that properly remain in the multicast group until they send join requests in response to the next general multicast listener query from the router.

To configure the router to immediately remove these groups, include the immediate-leave statement:

```
immediate-leave;
```

You can include this statement at the following hierarchy levels:

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

Configuring General IGMP Snooping Parameters

You can configure a number of general interval parameters that apply to all IGMP snooping interfaces. You can also configure these parameters to apply to a particular VLAN.

You can configure three intervals to determine the host-query message parameters used to time out a group:

- Query interval
- Query last-member interval

- Query response interval

To configure time-out group intervals, include the query-interval statement:

query-interval *seconds*;

By default, the query interval is 125 seconds. You can configure any value in the range 1 through 1024 seconds.

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 lower this interval to reduce the amount of time it takes to detect the loss of the last member of a group. To configure the maximum amount of time between group-specific query messages, include the query-last-member-interval statement:

query-last-member-interval *seconds*;

By default, the last-member query interval is 1 second. You can configure any value in the range 0.1 through 0.9 seconds, and then 1-second intervals from 1 through 1024 seconds.

To configure how long the router waits to receives a response from its host-query messages, include the query-response-interval statement:

query-response-interval *seconds*;

By default, the query response interval is 10 seconds. You can configure any value in the range 1 through 1024 seconds. However, this interval must be less than the interval set in the query-interval statement.

In addition to the intervals, the IGMP robustness variable provides fine-tuning to allow for expected packet loss on a subnet. It is basically the number of intervals to wait before timing out a group. You can wait more intervals if subnet packet loss is high and IGMP report messages might be lost. To configure a robustness variable, include the robust-count statement:

robust-count *number*;

By default, the robust count is 2. You can configure any value in the range 2 through 10 intervals.

You can configure general IGMP snooping statements at the following hierarchy levels:

- [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 igmp-snooping vlan *vlan-id*]

Configuring IGMP Snooping Interfaces

You can set many IGMP snooping parameters for a particular interface. Different interfaces can have different values of these parameters. To set the IGMP snooping parameters for a particular interface, include the **interface** statement:

```
interface interface-name {
   group-limit limit;
   host-only-interface;
   multicast-router-interface;
   static {
      group ip-address {
         source ip-address;
      }
   }
}
```

You can configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) that can join an interface. After this limit is reached, new reports will be ignored and all related flows will be discarded, not flooded. To configure the group limit for the interface, include the **group-limit** statement:

```
group-limit limit;
```

By default, there is no limit to the number of groups that can join an interface. You can configure a limit in the range 0 through a 32-bit number.

You can configure an IGMP snooping interface to be an exclusively router-facing or host-side interface. On a host-side interface, received IGMP queries are dropped. To configure a router-facing or host-facing interface, include either the **host-only-interface** or **multicast-router-interface** statement:

```
host-only-interface;
multicast-router-interface;
```

By default, an interface can face either other multicast routers or hosts.

You can configure an IGMP snooping interface with multicast groups statically, with or without source addresses. To configure the static groups for the interface, include the **static** statement:

```
static {
   group ip-address {
      source ip-address;
   }
}
```

By default, the router learns about multicast groups on the interface dynamically.

You can include this statement at the following hierarchy levels:

- [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 igmp-snooping vlan *vlan-id*]

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;
}
```

The operation of each statement is the same as described earlier in this chapter.

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]

Tracing IGMP Snooping Operations

To trace IGMP activity, you can include the traceoptions statement.

For general information about tracing and global tracing options, see the *JUNOS Routing Protocols Configuration Guide*.

Chapter 34

Summary of IGMP Snooping Configuration Statements

The following sections explain each of the IGMP snooping configuration statements. The statements are organized alphabetically.

group

Syntax

group *ip-address* {
  source-address *ip-address*;
}

Hierarchy Level

[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name* static],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name* static],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name* static],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name* static]

Release Information

Statement introduced in JUNOS Release 8.5.

Description

Configure the IGMP multicast group address that receives data on an interface and (optionally) a source address for certain packets.

Options

ip-address—Group address.

The remaining statement is explained separately.

Usage Guidelines

See “Configuring IGMP Snooping Interfaces” on page 297.

Required Privilege Level

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

group-limit

Syntax group-limit *limit*;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name*]

Release Information Statement introduced in JUNOS 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 will be 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 *limit*—a 32-bit number for the limit on the interface.

Usage Guidelines See “Configuring IGMP Snooping Interfaces” on page 297.

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

host-only-interface

Syntax host-only-interface;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name*]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring IGMP Snooping Interfaces" on page 297.

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

Related Topics
- multicast-router-interface

igmp-snooping

Syntax

```
igmp-snooping {
   immediate-leave;
   interface interface-name {
      group-limit limit;
      host-only-interface;
      immediate-leave;
      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;
   vlan vlan-id {
      immediate-leave;
      interface interface-name {
         group-limit limit;
         host-only-interface;
         immediate-leave;
         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;
   }
}
```

Hierarchy Level

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

Release Information

Statement introduced in JUNOS Release 8.5.

Description

Enable IGMP snooping on the router.

Default IGMP snooping is disabled on the router.

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 Topics
- *JUNOS Multicast Protocols Configuration Guide*

immediate-leave

Syntax immediate-leave;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[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 igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name*]

Release Information Statement introduced in JUNOS Release 8.5.

Description When this statement is enabled on a router running IGMPv2 snooping, after the router receives a leave group membership message from a host associated with the interface, the router immediately removes the group membership from the interface and suppresses the sending of any group-specific queries for the multicast group.

When this statement is enabled on a router running IGMPv3 snooping, after the router receives a report with the type BLOCK_OLD_SOURCES, the router suppresses the sending of group-and-source queries but relies on the JUNOS Software's host tracking mechanism to determine whether or not it removes a particular source group membership from the interface.

NOTE: When configuring this statement on IGMPv2 interfaces, ensure that the IGMP interface has only one IGMP host connected. If more than one IGMPv2 host is connected to a LAN through the same interface, and one host sends a leave message, the router removes all hosts on the interface from the multicast group. The router loses contact with the hosts that properly remain in the multicast group until they send join requests in response to the next general multicast listener query from the router.

Usage Guidelines See “Configuring General IGMP Snooping Parameters” on page 295.

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

interface

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 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. For details about specifying interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

The remaining statements are explained separately.

Usage Guidelines

See "Configuring IGMP Snooping Interfaces" on page 297.

Required Privilege Level

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

multicast-router-interface

Syntax multicast-router-interface;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name*]

Release Information Statement introduced in JUNOS Release 8.5.

Description Configure an interface as a bridge interface toward other multicast routers.

Default The interface can either be a host-side or multicast-router interface.

Usage Guidelines See “Configuring IGMP Snooping Interfaces” on page 297.

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

Related Topics host-only-interface

proxy

Syntax

```
proxy {
    source-address 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 Release 8.5.

Description

Configure proxy mode and options, including source address. All the queries generated by IGMP snooping are sent using 0.0.0.0 as source address in order to avoid participating in IGMP Querier election. Also, all reports generated by IGMP snooping are sent with 0.0.0.0 as source address unless there is a configured source address to use.

Default

By default, IGMP snooping does not employ proxy mode.

The remaining statement is explained separately.

Usage Guidelines

See “Configuring IGMP Snooping Proxy Mode” on page 294.

Required Privilege Level

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

query-interval

Syntax query-interval *seconds*;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*]

Release Information Statement introduced before JUNOS Release 8.5.

Description Configure the interval for host-query message timeouts.

Options *seconds*—Time interval.
Range: 1 through 1024
Default: 125 seconds

Usage Guidelines See "Configuring General IGMP Snooping Parameters" on page 295.

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

Related Topics
- query-response-interval
- query-last-member-interval

query-last-member-interval

Syntax query-last-member-interval *seconds*;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*]

Release Information Statement introduced in JUNOS Release 8.5.

Description Configure the interval for group-specific query timeouts.

Options *seconds*—Time interval, in fractions of a second or seconds.
Range: 0.1 through 0.9, then in 1-second intervals 1 through 1024
Default: 1 second

Usage Guidelines See "Configuring General IGMP Snooping Parameters" on page 295.

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

Related Topics
- query-response-interval
- query-interval

query-response-interval

Syntax query-response-interval *seconds*;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*]

Release Information Statement introduced in JUNOS Release 8.5.

Description Specify how long to wait to receive a response to a specific query message from a host.

Options *seconds*—Time interval. This interval must be less than the host-query interval.
Range: 1 through 1024
Default: 10 seconds

Usage Guidelines See "Configuring General IGMP Snooping Parameters" on page 295.

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

Related Topics
- query-last-member-interval
- query-interval

robust-count

Syntax robust-count *number*;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*]

Release Information Statement introduced in JUNOS Release 8.5.

Description Provide fine-tuning to allow for expected packet loss on a subnet. You can wait more intervals if subnet packet loss is high and IGMP report messages might be lost.

Options *number*—Robust interval.
Range: 2 through 10
Default: 2

Usage Guidelines See "Configuring General IGMP Snooping Parameters" on page 295.

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

source

Syntax source *ip-address*;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name* static group],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name* static group],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name* static group],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name* static group]

Release Information Statement introduced in JUNOS Release 8.5.

Description Statically define multicast group source addresses on an interface.

Options *ip-address*—IP address to use as source for group.

Usage Guidelines See "Configuring IGMP Snooping Interfaces" on page 297.

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

source-address

Syntax source-address *ip-address*;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping proxy],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* proxy],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping proxy],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* proxy]

Release Information Statement introduced in JUNOS Release 8.5.

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.

Options *ip-address*—IP address to use as the source for proxy-mode IGMP snooping reports.

Usage Guidelines See “Configuring IGMP Snooping Proxy Mode” on page 294.

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

static

Syntax

```
static {
   group multicast-group-address {
      source ip-address;
   }
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*]

Release Information

Statement introduced in JUNOS Release 8.5.

Description

Define static multicast groups on an interface.

The remaining statements are explained separately.

Usage Guidelines

See “Configuring IGMP Snooping Interfaces” on page 297.

Required Privilege Level

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

vlan

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced in JUNOS Release 8.5.

Description

Configure IGMP snooping parameters for a particular VLAN.

Default

By default, IGMP snooping options apply to all VLANs.

Options

vlan-id—Apply the parameters to this VLAN.

The remaining statements are explained separately.

Usage Guidelines

See “Configuring VLAN-Specific IGMP Snooping Parameters” on page 298.

Required Privilege Level

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

Related Topics

igmp-snooping

Part 10
DVMRP

- DVMRP Overview on page 317
- DVMRP Configuration Guidelines on page 319
- Summary of DVMRP Configuration Statements on page 329

Chapter 35

DVMRP Overview

The Distance Vector Multicast Routing Protocol (DVMRP) is a distance-vector routing protocol that provides connectionless datagram delivery to a group of hosts across an internetwork. DVMRP is a distributed protocol that dynamically generates IP multicast delivery trees by using a technique called reverse-path multicasting (RPM) to forward multicast traffic to downstream interfaces. These mechanisms allow the formation of shortest-path trees, which are used to reach all group members from each network source of multicast traffic.

DVMRP is designed to be used as an interior gateway protocol (IGP) within a multicast domain.

Because not all IP routers support native multicast routing, DVMRP includes direct support for tunneling IP multicast datagrams through routers. The IP multicast datagrams are encapsulated in unicast IP packets and addressed to the routers that do support native multicast routing. DVMRP treats tunnel interfaces and physical network interfaces the same way.

DVMRP routers dynamically discover their neighbors by sending neighbor probe messages periodically to an IP multicast group address that is reserved for all DVMRP routers.

For information about standards supported for DVMRP, see “IP Multicast Specifications” on page 18.

Chapter 36

DVMRP Configuration Guidelines

This chapter describes the following tasks for configuring DVMRP:

- Configuring DVMRP on page 319
- Minimum DVMRP Configuration on page 320
- Creating Routing Tables for DVMRP Routes on page 321
- Enabling DVMRP on page 321
- Modifying the DVMRP Hold-Time Period on page 322
- Modifying the DVMRP Metric Value on page 322
- Disabling DVMRP on an Interface on page 323
- Configuring DVMRP Routing Policy on page 323
- Configuring DVMRP Routing Modes on page 324
- Tracing DVMRP Protocol Traffic on page 324
- DVMRP Configuration Examples on page 325
- Example: Tracing DVMRP Protocol Traffic on page 325
- Example: Configuring DVMRP on page 326
- Example: Configuring DVMRP to Announce Unicast Routes on page 326

Configuring DVMRP

To configure the Distance Vector Multicast Routing Protocol (DVMRP), include the dvmrp statement:

```
dvmrp {
    disable;
    export [ policy-names ];
    import [ policy-names ];
    interface interface-name {
        disable;
        hold-time seconds;
        metric metric;
        mode (forwarding | unicast-routing);
    }
    rib-group group-name;
    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]

For an overview of logical systems and a detailed example of logical system configuration, see the logical systems chapter of the *JUNOS Feature Guide*.

By default, DVMRP is disabled.

Minimum DVMRP Configuration

To enable DVMRP on an interface, include at least the following statements in the configuration. All other DVMRP configuration statements are optional.

```
routing-options {
    interface-routes {
        rib-group group-name1;
    }
    rib-groups {
        group-name1 {
            import-rib [ inet.0 inet.2 ];
        }
        group-name2 {
            import-rib inet.2;
            export-rib inet.2;
        }
    }
}
protocols {
    dvmrp {
        rib-group group-name2;
        interface interface-name;
    }
}
```

You can include these statements at the following hierarchy levels:

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

The port of a DVMRP router can be either a physical interface to a directly attached subnetwork or a tunnel interface to another multicast-capable area of the Multicast Backbone (MBone). All interfaces can be configured with a metric specifying cost for receiving packets on a given port. The default metric is 1.

Creating Routing Tables for DVMRP Routes

DVMRP needs to access route information from the unicast routing table, inet.0, and from a separate routing table that is reserved for DVMRP. You need to create the routing table for DVMRP and to create groups of routing tables so that the routing protocol process imports and exports routes properly. We recommend that you use routing table inet.2 for DVMRP routing information.

To create the necessary routing tables and routing table groups for DVMRP, include the following statements:

```
routing-options {
    interface-routes {
        rib-group group-name1;
    }
    rib-groups {
        group-name1 {
            import-rib [ inet.0 inet.2 ];
        }
        group-name2 {
            import-rib inet.2;
            export-rib inet.2;
        }
    }
}
```

You can include these statements at the following hierarchy levels:

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

To associate the routing tables with DVMRP, include the rib-group statement at the [edit protocols dvmrp] hierarchy level, as described in "Enabling DVMRP" on page 321.

Enabling DVMRP

To enable DVMRP on the router, include the following statements:

```
dvmrp {
    interface interface-name;
    rib-group group-name;
}
```

You can include these statements at the following hierarchy levels:

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

The rib-group statement selects a routing table group. DVMRP exports routes from this group and imports routes to this group. The rib-group statement associates with DVMRP the routing table group that imports and exports routes into the specified routing table group. This is a group you defined with the rib-groups statement at the [edit routing-options] hierarchy level.

You must specify the interface or interfaces on which to enable DVMRP. Specify the full interface name, including the physical and logical address components. To configure all interfaces, specify the interface name all. For details about specifying interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

NOTE: If you have configured Protocol Independent Multicast (PIM) on the interface, you can configure DVMRP in unicast-routing mode only. You cannot configure PIM and DVMRP in forwarding mode at the same time.

Modifying the DVMRP Hold-Time Period

The DVMRP hold-time period is the amount of time that a neighbor is to consider the sending router (this router) to be operative (up). The default hold-time period is 35 seconds.

To modify the hold-time value for the local router, include the hold-time statement:

hold-time *seconds*;

You can include this statement at the following hierarchy levels:

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

The hold-time period can range from 1 through 255 seconds.

Modifying the DVMRP Metric Value

For each source network reported, a route metric is associated with the unicast route being reported. The metric is the sum of the interface metrics between the router originating the report and the source network. A metric of 32 marks the source network as unreachable, thus limiting the breadth of the DVMRP network and placing an upper bound on the DVMRP convergence time.

By default, a metric value of 1 is associated with each DVMRP route. To modify the metric value, include the metric statement:

metric *metric*;

You can include this statement at the following hierarchy levels:

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

The metric can range from 1 through 31.

Disabling DVMRP on an Interface

To disable DVMRP on an interface, include the disable statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

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

Configuring DVMRP Routing Policy

All routing protocols use the routing table to store the routes that they learn and to determine which routes they advertise in their protocol packets. Routing policy allows you to control which routes the routing protocols store in and retrieve from the routing table.

When configuring DVMRP routing policy, you can apply routing policies. To do this, include the import and export statements at the [edit protocols dvmrp] or [edit logical-systems *logical-system-name* protocols dvmrp] hierarchy level.

To apply policies to routes imported into the routing table from DVMRP, include the import statement, listing the names of one or more policy filters to be evaluated. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found, DVMRP shares with the routing table only those routes that were learned from DVMRP routers.

```
import [ policy-names ];
```

You can include this statement at the following hierarchy levels:

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

To apply policies to routes exported from the routing table into DVMRP, include the export statement, listing the names of one or more policies to be evaluated. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found, the routing table exports into DVMRP only the routes that it learned from DVMRP and direct routes.

```
export [ policy-names ];
```

You can include this statement at the following hierarchy levels:

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

Configuring DVMRP Routing Modes

You can configure DVMRP for either forwarding or unicast routing mode. In forwarding mode, DVMRP operates its protocol normally (for example, it does the routing as well as multicast data forwarding). In unicast routing mode, you can use DVMRP for unicast routing only; to forward multicast data, enable PIM on that interface. To configure the mode, include the **mode** statement.

To configure DVMRP for multicast forwarding, include the **mode forwarding** statement:

```
mode forwarding;
```

To configure DVMRP for unicast routing, include the **mode unicast-routing** statement:

```
mode unicast-routing;
```

You can include these statements at the following hierarchy levels:

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

The default mode is forwarding.

Tracing DVMRP Protocol Traffic

To trace DVMRP protocol traffic, you can specify options in the global **traceoptions** statement at the [edit routing-options] or [edit logical-systems *logical-system-name* routing-options] hierarchy level. Options applied at the routing options level trace all packets, and options applied at the protocol level trace only DVMRP traffic.

You can specify DVMRP-specific options by including the **traceoptions** statement:

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

You can specify the following DVMRP-specific options in the DVMRP **traceoptions** statement:

- **all**—Trace everything.
- **general**—Trace general events.
- **graft**—Trace graft messages.
- **neighbor**—Trace neighbor probe messages.

- normal—Trace normal events.
- packets—Trace all DVMRP packets.
- poison—Trace poison-route-reverse packets.
- policy—Trace policy processing.
- probe—Trace probe packets.
- prune—Trace prune messages.
- report—Trace DVMRP route report packets.
- route—Trace routing information.
- state—Trace state transitions.
- task—Trace routing protocol task processing.
- timer—Trace routing protocol timer processing.

For general information about tracing and global tracing options, see the *JUNOS Routing Protocols Configuration Guide*.

DVMRP Configuration Examples

This section contains the following DVMRP configuration examples:

- Example: Tracing DVMRP Protocol Traffic on page 325
- Example: Configuring DVMRP on page 326
- Example: Configuring DVMRP to Announce Unicast Routes on page 326

Example: Tracing DVMRP Protocol Traffic

Trace only unusual or abnormal operations to the file routing-log, and trace detailed information about all DVMRP messages to the file dvmrp-log:

```
[edit]
routing-options {
   traceoptions {
      file routing-log;
   }
}
protocols {
   dvmrp {
      traceoptions {
         file dvmrp-log;
         flag packets;
      }
      interface so-0/0/0;
   }
}
```

Example: Configuring DVMRP

Configure DVMRP on the router:

```
[edit]
routing-options {
   interface-routes {
      rib-group ifrg;
   }
   rib-groups {
      ifrg {
         import-rib [ inet.0 inet.2 ];
      }
      dvmrp-rib {
         import-rib inet.2;
         export-rib inet.2;
      }
   }
}
protocols {
   sap;
   dvmrp {
      rib-group dvmrp-rib;
      traceoptions {
         flag normal;
         flag state;
      }
      interface ip-f/p/0.0 {
         hold-time 130;
      }
   }
}
```

Example: Configuring DVMRP to Announce Unicast Routes

In this example, DVMRP is used to announce unicast routes used solely for multicast reverse-path forwarding (RPF). Include the **mode unicast-routing** statement at the **[edit protocols dvmrp interface]** hierarchy level. Redistribute static routes by including the **static** statement at the **[edit routing-options]** hierarchy level to export the routes to all DVMRP neighbors.

```
[edit]
routing-options {
   rib inet.2 {
      static {
         route 0.0.0.0/0 discard;
      }
   }
   rib-groups {
      pim-rg {
         import-rib inet.2;
      }
      dvmrp-rg {
```

```
            export-rib inet.2;
            import-rib inet.2;
        }
    }
protocols {
    dvmrp {
        rib-group inet dvmrp-rg;
    export dvmrp-export;
        interface all {
            mode unicast-routing;
        }
    pim {
        rib-group inet pim-rg;
        interface all;
    }
}
policy-options {
    policy-statement dvmrp-export {
        term 10 {
            from {
                protocol static;
                route-filter 0.0.0.0/0 exact;
            }
            then accept;
        }
    }
}
```

Chapter 37

Summary of DVMRP Configuration Statements

The following sections explain each of the Distance Vector Multicast Routing Protocol (DVMRP) configuration statements. The statements are organized alphabetically.

disable

Syntax disable;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Explicitly disable DVMRP on the system or on an interface.

Usage Guidelines See "Disabling DVMRP on an Interface" on page 323.

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

dvmrp

Syntax

```
dvmrp {
   disable;
   export [ policy-names ];
   import [ policy-names ];
   interface interface-name {
      disable;
      hold-time seconds;
      metric metric;
      mode (forwarding | unicast-routing);
   }
   rib-group group-name;
   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 Release 7.4.

Description

Enable DVMRP on the router.

Default

DVMRP is disabled on the router.

Options

The statements are explained separately.

Usage Guidelines

See "Enabling DVMRP" on page 321.

Required Privilege Level

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

export

Syntax export [*policy-names*];

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Apply one or more policies to routes being exported from the routing table into DVMRP.

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

Usage Guidelines See "Configuring DVMRP Routing Policy" on page 323.

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

Related Topics import

hold-time

Syntax hold-time *seconds*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the time period for which a neighbor is to consider the sending router (this router) to be operative (up).

Options *seconds*—Hold time.
Range: 1 through 255
Default: 35 seconds

Usage Guidelines See "Modifying the DVMRP Hold-Time Period" on page 322.

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

import

Syntax import [*policy-names*];

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Apply one or more policies to routes being imported into the routing table from DVMRP.

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

Usage Guidelines See "Configuring DVMRP Routing Policy" on page 323.

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

Related Topics export

interface

Syntax

```
interface interface-name {
   disable;
   hold-time seconds;
   metric metric;
   mode (forwarding | unicast-routing);
}
```

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Enable DVMRP 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. For details about specifying interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

The remaining statements are explained separately.

Usage Guidelines See "Enabling DVMRP" on page 321.

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

metric

Syntax metric *metric*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the DVMRP metric value.

Options *metric*—Metric value.
Range: 1 through 31
Default: 1

Usage Guidelines See "Modifying the DVMRP Metric Value" on page 322.

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

mode

Syntax mode (forwarding | unicast-routing);

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure DVMRP for multicast traffic forwarding or unicast routing.

Options forwarding—DVMRP performs unicast routing as well as multicast data forwarding.

unicast-routing—DVMRP performs unicast routing only. To forward multicast data, you must configure Protocol Independent Multicast (PIM) on the interface.

Usage Guidelines See "Configuring DVMRP Routing Modes" on page 324.

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

rib-group

Syntax rib-group *group-name*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Associate a routing table group with DVMRP.

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.

Usage Guidelines See "Enabling DVMRP" on page 321.

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

traceoptions

Syntax

```
traceoptions {
   file filename <files number> <size size> <world-readable | no-world-readable>;
   flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Configure DVMRP tracing options.

To specify more than one tracing operation, include multiple flag statements.

Default

The default DVMRP trace options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place tracing output in the dvmrp-log file.

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.
Range: 2 through 1000 files
Default: 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.

DVMRP Tracing Flags

- all—All tracing operations
- general—A combination of the normal and route trace operations
- graft—Graft messages
- neighbor—Neighbor probe messages
- normal—All normal operations

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

- packets—All DVMRP packets
- poison—Poison-route-reverse packets
- probe—Probe packets
- prune—Prune messages
- report—DVMRP route report packets
- policy—Policy operations and actions
- route—Routing table changes
- state—State transitions
- task—Interface transactions and processing
- timer—Timer usage

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

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

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.
Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

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

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

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

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

Syntax: *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB

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

Default: 1 MB

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

Usage Guidelines

See "Tracing DVMRP Protocol Traffic" on page 324.

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.

Part 11
PIM

- PIM Overview on page 341
- PIM Mode-Independent Configuration Guidelines on page 369
- PIM Dense and Sparse-Dense Mode Configuration Guidelines on page 391
- PIM Sparse Mode Configuration Guidelines on page 393
- PIM Configuration Examples on page 421
- Summary of PIM Configuration Statements on page 445

Chapter 38

PIM Overview

For more information about how the PIM modes operate, see the following sections:

- Protocol Independent Multicast Overview on page 341
- PIM Sparse Mode on page 342
- Designated Router on page 344
- Rendezvous Point on page 344
- RP Mapping Options on page 344
- Building an RPT Between RP and Receivers on page 346
- PIM Sparse-Mode Source Registration on page 347
- PIM Sparse-Mode SPT Cutover on page 350
- SPT Cutover on page 350
- SPT Cutover Control on page 354
- PIM SSM on page 354
- PIM Dense Mode on page 355
- PIM Sparse-Dense Mode on page 357
- RP Mapping with Anycast RP on page 357
- Multicast over Layer 3 VPNs on page 358
- Dual PIM Multicast VPNs on page 358
- MBGP-Based Multicast VPNs on page 359
- Tunnel Services PICs and Multicast on page 362
- Filtering Multicast Messages on page 363
- Filtering MAC Addresses on page 364
- Filtering RP/DR Register Messages on page 364
- Filtering MSDP SA Messages on page 365
- Embedded RP for IPv6 Multicast on page 366

Protocol Independent Multicast Overview

Protocol Independent Multicast (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. The JUNOS Software supports sparse mode, dense mode, and sparse-dense mode.

For information about standards supported for PIM, see "IP Multicast Specifications" on page 18.

Because the PIM mode you choose determines the PIM configuration properties, you first must decide whether PIM operates in 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 router unless it 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 joins and prunes are common.

Unlike sparse 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.

PIM Sparse Mode

A 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 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's 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's DR encapsulates the packets in a PIM register message and forwards it 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 simply 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 gets the first packet from the RPT, the DR sends a PIM join message toward the source's DR to start building an SPT back to the source. When the source's 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 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 knows of the active source's existence and no longer needs to receive advertisement of the source (which utilizes resources).

This section contains more information about the routers and PIM sparse-mode functions briefly described above:

- Designated Router on page 344
- Rendezvous Point on page 344
- RP Mapping Options on page 344
- Building an RPT Between RP and Receivers on page 346
- PIM Sparse-Mode Source Registration on page 347
- PIM SSM on page 354

Designated Router

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

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

Regardless of whether it is the receiver's DR or the source's DR, a DR is selected from other routers in a network by the exchange of IP addresses. Neighboring PIM sparse-mode 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 Software, 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.

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 get to 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 SPT. As shown in Figure 15 on page 344, the RP router is upstream from the receiver and thus forms one end of the RPT.

Figure 15: Rendezvous Point as Part of the RPT and SPT

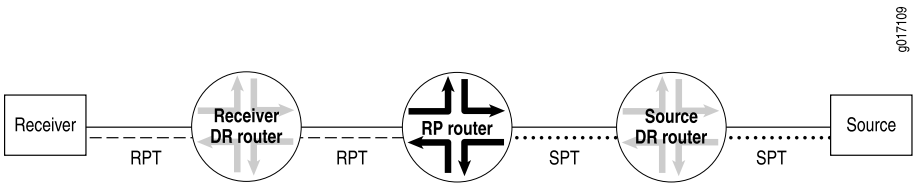

The benefit of using the RP as the information exchange point is that it reduces the amount of state in non-RP routers. No network flooding is required to provide non-RP routers information about active sources.

RP Mapping Options

RPs can be learned by one of the following mechanisms.

Static Configuration

You can configure a static RP configuration that is very 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.

Anycast RP

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.

Having a single active 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.

When anycast RP is configured, the shared address is used in the RP-to-group mapping. This allows multicast groups to have multiple active RPs in a PIM domain. However, the RPs must use some protocol to synchronize the active source information so that the active RP for each group is known to all RPs.

There are two methods for RP active source synchronization in anycast RP, one using the Multicast Source Discovery Protocol (MSDP) and the other using PIM itself.

When MSDP is used with PIM sparse mode, anycast RP provides a faster failover rate than auto-RP or a bootstrap router. However, MSDP only works for IPv4. When PIM alone is used for anycast RP, the solution works for both IPv4 and IPv6.

For more information about configuring static RPs, see “Configuring Static PIM RPs” on page 399. For more information about configuring anycast RP, see “Configuring PIM Auto-RP” on page 402 and “Example: Configuring PIM Anycast with MSDP” on page 423.

Auto-RP

You can configure a more dynamic way of assigning 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 stops operating, one of the other preconfigured routers takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

For more information, see “Configuring PIM Auto-RP” on page 402.

Bootstrap Router

To determine which router is the 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's bootstrap router originates 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.

For more information, see “Configuring PIM Bootstrap Properties” on page 400.

Building an RPT Between RP and Receivers

The RPT is the path between the RP and receivers (hosts) in a multicast group (see Figure 16 on page 347). 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 16: Building an RPT Between RP and Receiver

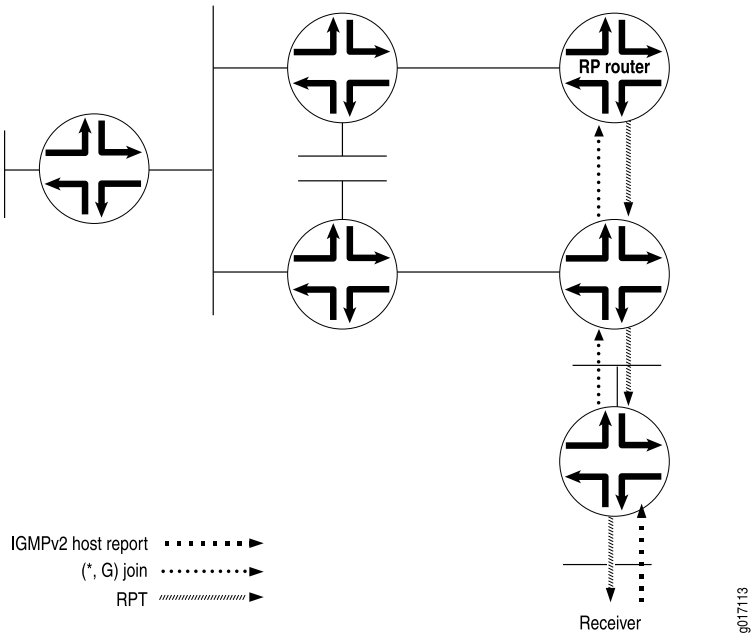

PIM Sparse-Mode Source Registration

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 SPT needs to be built from the source's DR to the RP.

The SPT 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 out to the RP router (see Figure 17 on page 348).
2. When the RP router receives the PIM register message from the source, it sends a PIM join message back to the source.

Figure 17: PIM Register Message and PIM Join Message Exchanged

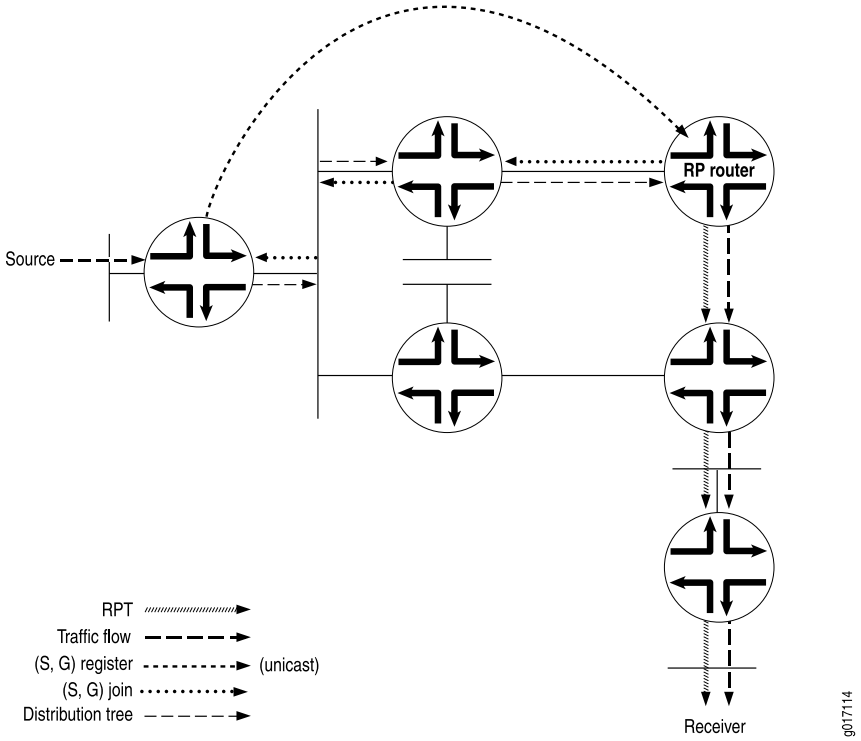

3. The source's DR receives the PIM join message and begins sending traffic down the SPT toward the RP router (see Figure 18 on page 349).
4. Once traffic is received by the RP router, it sends a register stop message to the source's DR to stop the register process.

Figure 18: Traffic Sent from the Source to the RP Router

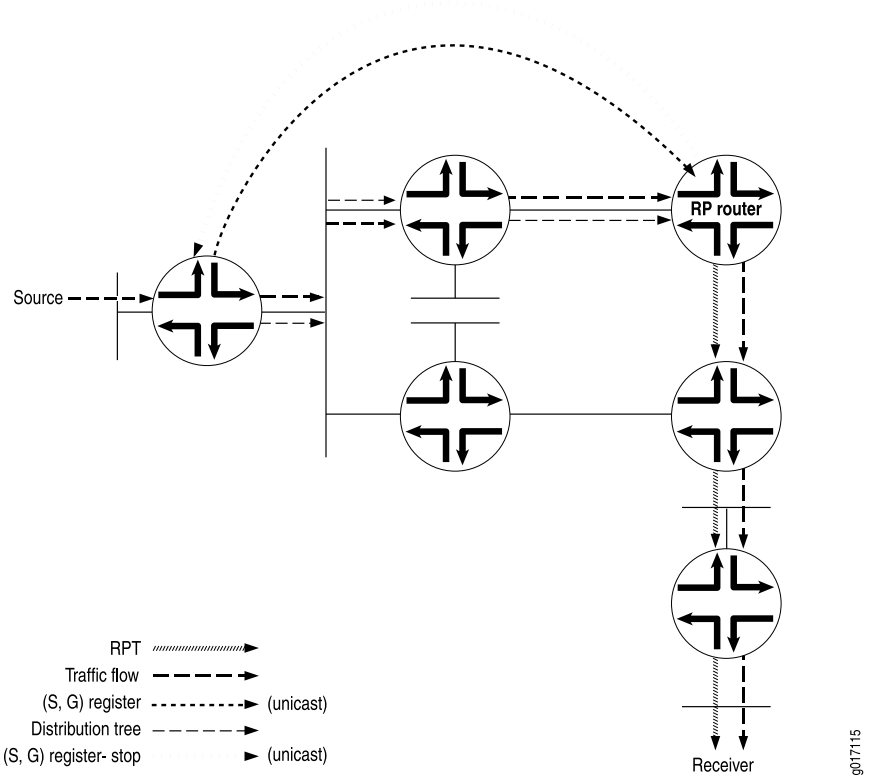

5. The RP router sends the multicast traffic down the RPT toward the receiver (see Figure 19 on page 350).

Figure 19: Traffic Sent from the RP Router Toward the Receiver

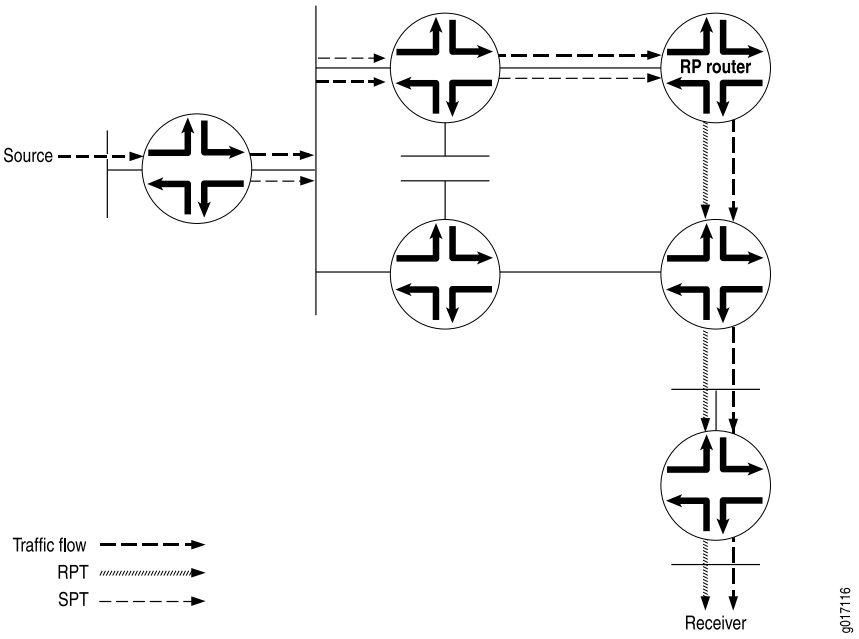

PIM Sparse-Mode SPT Cutover

The RPT is not always the most direct path for delivering multicast traffic to a receiver. In many cases, a direct SPT from the last-hop router to the source is a better way to receive a multicast stream. For more information, see the following sections:

- SPT Cutover on page 350
- SPT Cutover Control on page 354

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 20 on page 351).
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 20: Receiver DR Sends a PIM Join Message to the Source

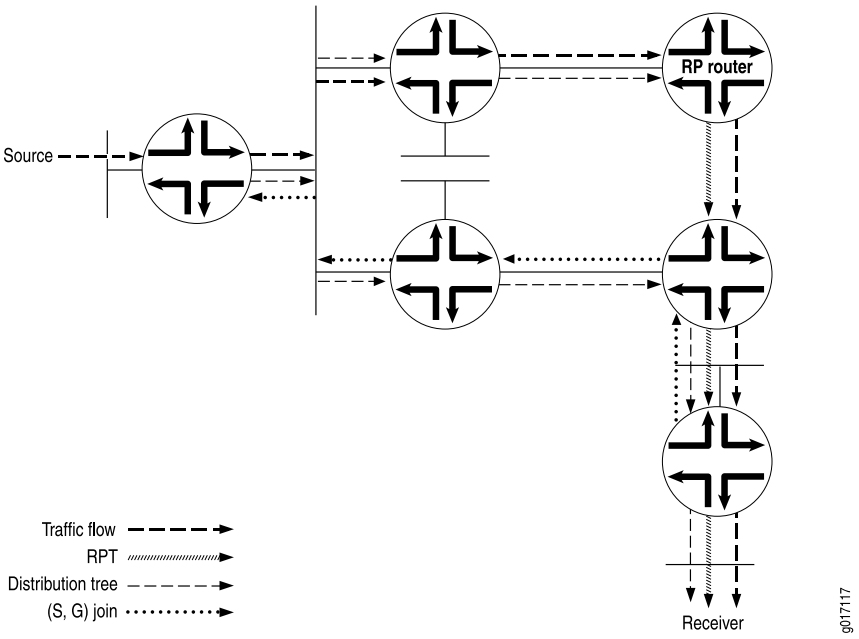

4. To stop duplicate multicast packets, the receiver's DR sends a PIM prune message toward the RP router, letting it know that the multicast packets from this particular source coming in from the RPT are no longer needed (see Figure 21 on page 351).

Figure 21: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router

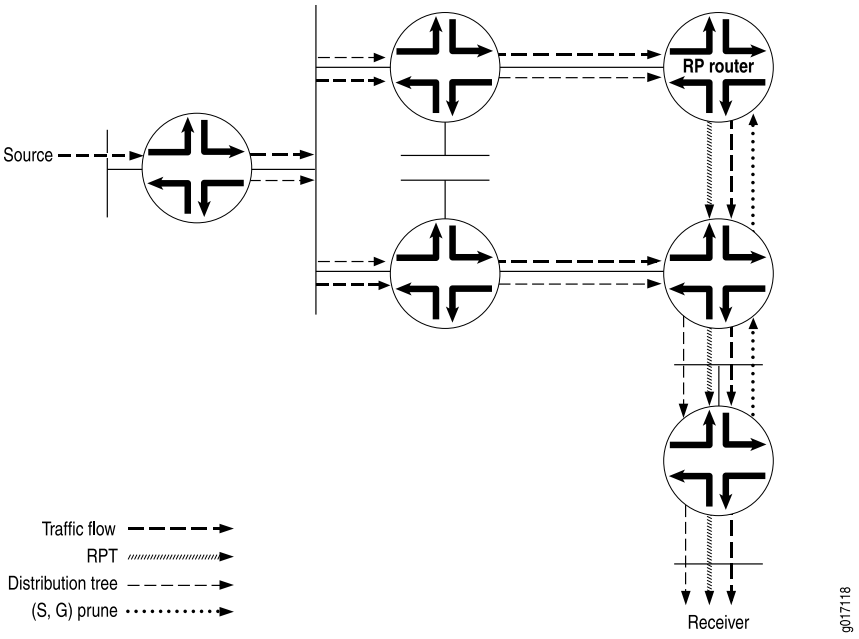

5. The PIM prune message is received by the RP router, and it stops sending multicast packets down to the receiver's DR. The receiver's DR is getting multicast packets only for this particular source over the new SPT. However, multicast packets from the source are still arriving from the source's DR toward the RP router (see Figure 22 on page 352).

Figure 22: RP Router Receives PIM Prune Message

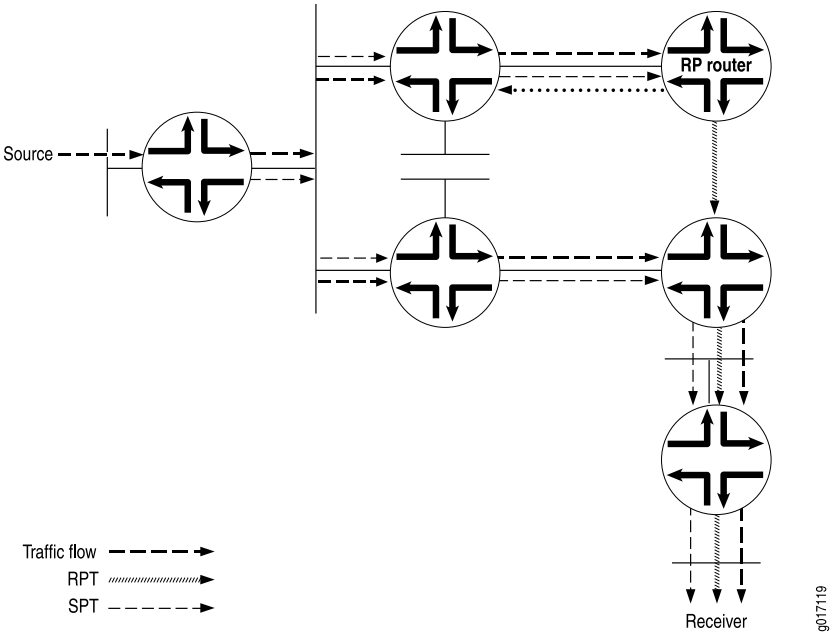

6. To stop the unneeded multicast packets from this particular source, the RP router sends a PIM prune message to the source's DR (see Figure 23 on page 353).

Figure 23: RP Router Sends a PIM Prune Message to the Source DR

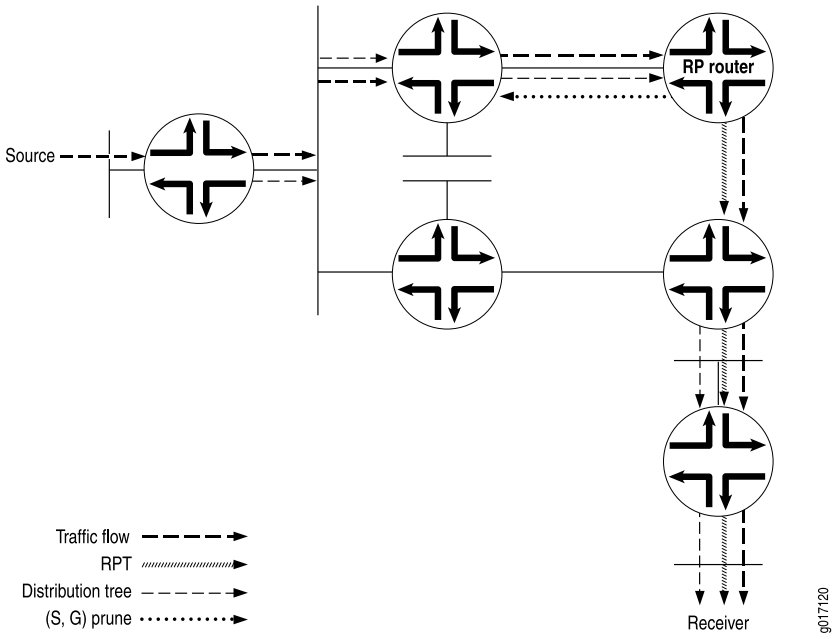

7. The receiver's DR now receives multicast packets only for the particular source from the SPT (see Figure 24 on page 353).

Figure 24: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router

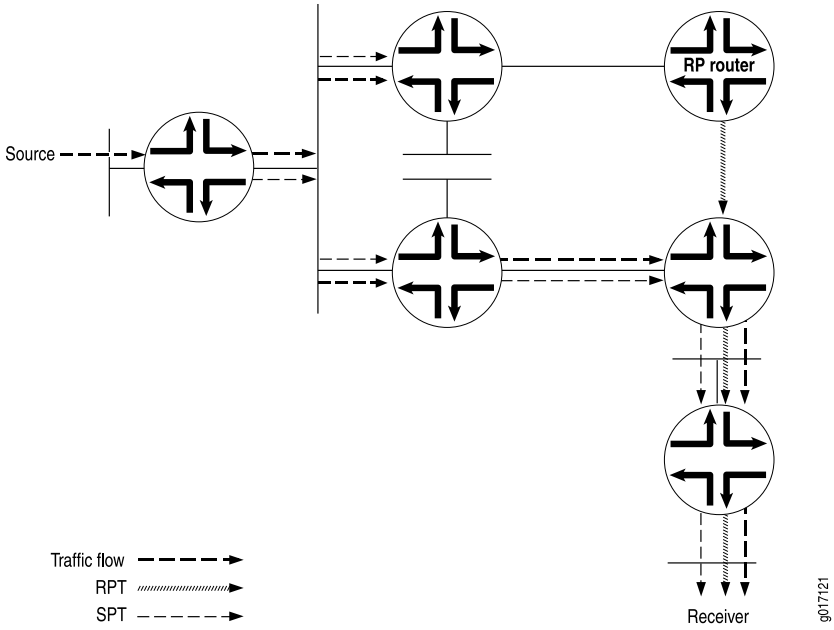

SPT Cutover Control

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.

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 address statement at the [edit routing-options multicast ssm-groups] hierarchy level.

You can also configure the JUNOS Software 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 only 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 355). 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 355 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

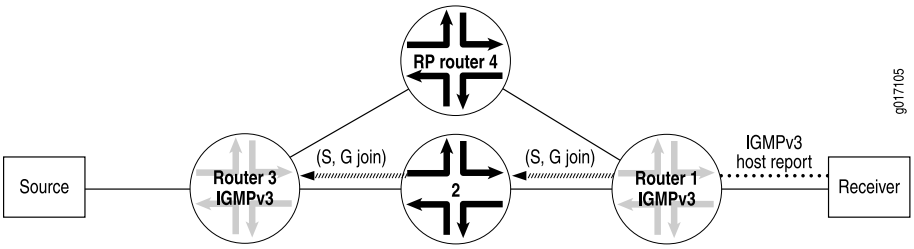

The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In Figure 26 on page 355, the source tree is built across the network to Router 3, the last-hop router connected to the source.

Figure 26: Router 3 (Last-Hop Router) Joins the Source Tree

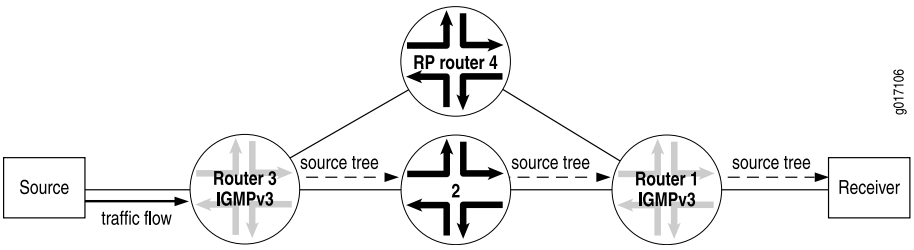

Using the source tree, multicast traffic is delivered to the subscribing host (see Figure 27 on page 355).

Figure 27: The (S,G) State Is Built Between the Source and the Receiver

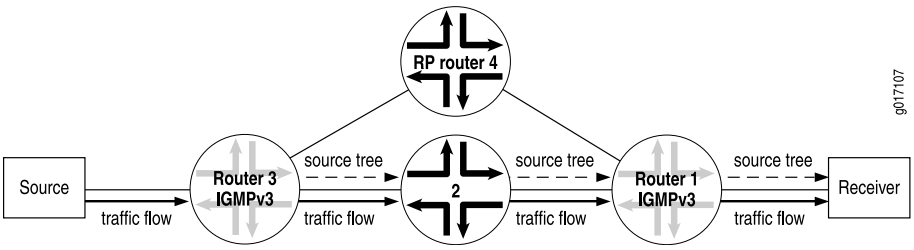

To configure additional SSM groups, include the ssm-groups statement at the [edit routing-options multicast] hierarchy level.

For more information about PIM SSM, see “Example: Configuring PIM SSM on a Network” on page 206.

PIM Dense Mode

Unlike sparse mode, in which data is forwarded only to routers sending an explicit request, dense mode implements a *flood-and-prune* mechanism, similar to DVMRP. In PIM dense mode, there is no RP. A router receives the multicast data on the

interface closest to the source, then forwards the traffic to all other interfaces (see Figure 28 on page 356).

Figure 28: Multicast Traffic Flooded from the Source Using PIM Dense Mode

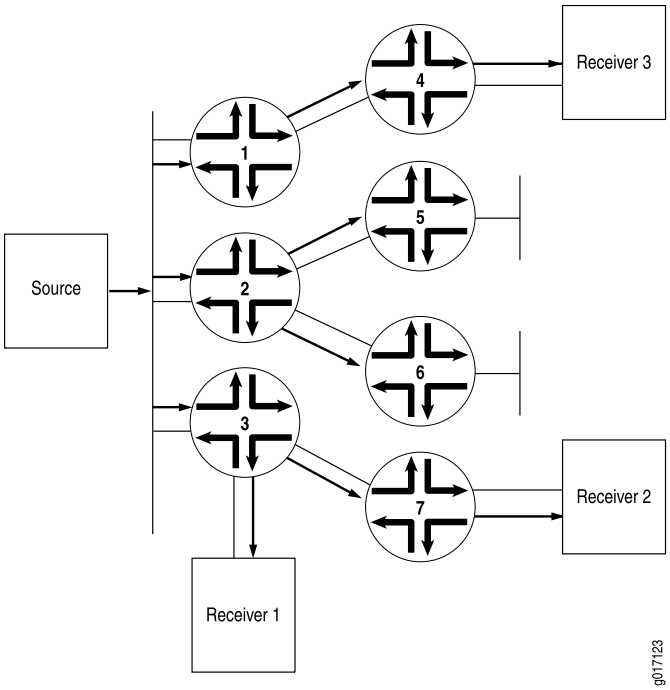

Flooding occurs periodically. It is used to refresh state information, such as the source IP address and multicast group pair. If the router has no interested receivers for the data, and the OIL becomes empty, the router sends a prune message upstream to stop delivery of multicast traffic (see Figure 29 on page 357).

Figure 29: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic

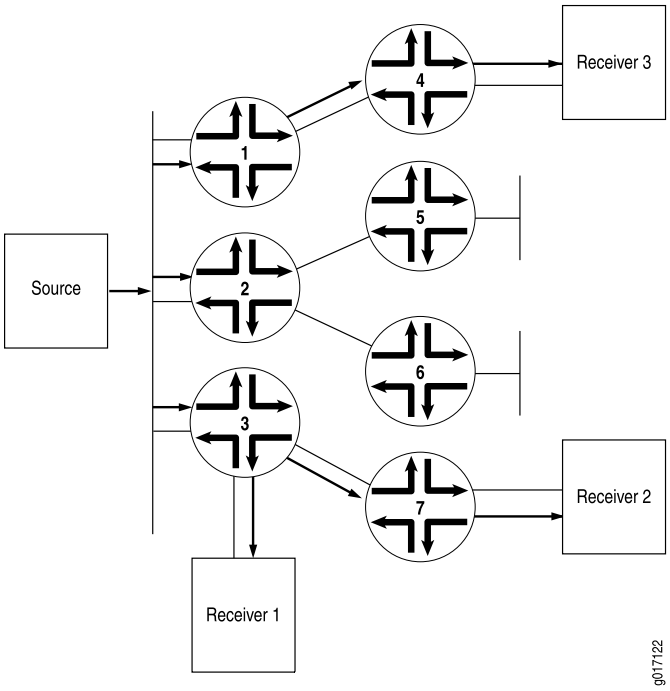

PIM Sparse-Dense Mode

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 “PIM Sparse Mode” on page 342 and “PIM Dense Mode” on page 355.

RP Mapping with Anycast RP

For the purposes of load balancing and redundancy, you can configure anycast RP. 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 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 goes down, sources and receivers are taken to a new RP by means of unicast routing.

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 Web site at http://www.ietf.org.

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.

See also "Example: Configuring PIM Anycast with MSDP" on page 423.

Multicast over Layer 3 VPNs

There are two ways to implement Layer 3 VPNs using PIM, or multicast VPNs (MVPNs). There are no official names for the two methods: using *dual PIM MVPNs* (also known informally as "draft-rosen") and *multiprotocol BGP (MBGP)-based MVPNs* (the "next generation" method of MVPN configuration). Both methods are supported and equally effective, but the MBGP-based MVPN method does not require multicast configuration on the service provider backbone. In other words, the PIM state information is maintained between the PE routers using the same architecture used for unicast VPNs. The main advantage of deploying MVPNs with MBGP is simplicity of configuration and operation because multicast is not needed on the service provider VPN backbone connecting the PE routers.

This section discusses two types of MVPNs:

- Dual PIM Multicast VPNs on page 358
- MBGP-Based Multicast VPNs on page 359

Dual PIM Multicast VPNs

In the unicast environment for Layer 3 virtual private networks (VPNs), all VPN state information is contained within the provider edge (PE) routers. However, with multicast for Layer 3 VPNs, PIM adjacencies are established in one of the following ways:

- You can set PIM adjacencies between the customer edge (CE) router and the PE router through a VPN routing and forwarding (VRF) instance at the [edit routing-instances *instance-name* protocols pim] hierarchy level. You must include the vpn-group-address statement at this hierarchy level, specifying a multicast group. The RP listed in the VRF-instance is the VPN customer RP (C-RP).
- You can also set the master PIM instance and the PE's IGP neighbors by configuring statements at the [edit protocols pim] hierarchy level. You must add the multicast group specified in the VRF instance to the master PIM instance. The set of master PIM adjacencies throughout the service provider network makes up the forwarding path that becomes an RP tree rooted at the service provider RP (SP-RP). Therefore, provider (P) routers within the provider core must maintain multicast state information for the VPNs.

For this configuration to work properly, you need two types of RP routers for each VPN:

- A VPN C-RP—An RP router located somewhere within the customer VPN
- An SP-RP—An RP router located within the service provider network

NOTE: A PE router can act as an SP-RP or the VPN C-RP of a VPN. However, when auto-RP and BSR are used, the PE cannot be a C-RP. It can, however, learn another router as C-RP by means of the auto-RP or BSR protocols.

For more information about configuring multicast for Layer 3 VPNs using dual PIM, see "Configuring Any-Source Multicast for Draft-Rosen VPNs" on page 415. For multicast Layer 3 VPN examples, see "Example: Configuring PIM Sparse Mode over Layer 3 VPNs" on page 436.

MBGP-Based Multicast VPNs

MBGP-based MVPNs (next generation MVPNs) are based on Internet drafts and extend unicast VPNs based on RFC 2547 to include support for IP multicast traffic. These MVPNs follow the same architectural model as the unicast VPNs and use BGP as the PE-to-PE control plane to exchange information. The next generation MVPN approach is based on Internet drafts draft-ietf-l3vpn-2547bis-mcast.txt, draft-ietf-l3vpn-2547bis-mcast-bgp.txt, and draft-morin-l3vpn-mvpn-considerations.txt.

In addition, MBGP-based MVPNs:

- Do not require a virtual router as in dual PIM MVPNs.
- Extend RFC 2547 unicast VPN mechanisms for intra-AS and extra-AS environments.
- Retain the scalability and flexibility of RFC 2547 unicast VPNs as much as possible.

MBGP-based MVPNs introduce two new types of tree:

Inclusive tree—A single multicast distribution tree in the backbone carrying all the multicast traffic from a specified set of one or more MVPNs. An inclusive tree carrying the traffic of more than one MVPN is an *aggregate inclusive tree*. All the PEs that attach to MVPN receiver sites using the tree belong to that inclusive tree.

Selective tree—A single multicast distribution tree in the backbone carrying traffic for a specified set of one or more multicast groups. When multicast groups belonging to more than one MVPN are on the tree, it is called an aggregate selective tree.

By default, traffic from most multicast groups can be carried by an inclusive tree, while traffic from some groups (for example, high bandwidth groups) can be carried by one of the selective trees. Selective trees, if they contain only those PEs that need to receive multicast data from one or more groups assigned to the tree, can provide

more optimal routing than inclusive tress alone, although this requires more state information in the P routers.

An MPLS-based VPN running BGP with autodiscovery is used as the basis for a next-generation MVPN. The autodiscovered route information is carried in MBGP network layer reachability information (NLRI) updates for multicast VPNs (MCAST-VPNs). These MCAST-VPN NLRIs are handled in the same way as IPv4 routes: route distinguishers are used to distinguish between different VPNs in the network. These NLRIs are imported and exported based on the route target extended communities, just as IPv4 unicast routes. In other words, existing BGP mechanisms are used to distribute multicast information on the provider backbone without requiring multicast directly.

For example, consider a customer running PIM sparse mode in SSM mode. Only source tree join customer multicast (c-multicast) routes are required. (PIM sparse mode in ASM mode can be supported with a few enhancements to SSM mode.)

The customer multicast route carrying a particular multicast source S needs to be imported only into the VRF table on the PE router connected to the site that contains the source S and not into any other VRF, even for the same MVPN. To do this, each VRF on a particular PE has a distinct VRF route import extended community associated with it. This community consists of the PE router's IP address and local PE number. Different MVPNs on a particular PE have different route imports, and for a particular MVPN the VRFs on different PE routers have different route imports. This VRF route import is auto-configured and not controlled by the user.

Also, all the VRFs within a particular MVPN will have information about VRF route imports for each VRF. This is accomplished by “piggybacking” the VRF route import extended community onto the unicast VPN IPv4 routes. In order to make sure a customer multicast route carrying multicast source S is imported only into the VRF on the PE router connected to the site contained the source S, it is necessary to find the unicast VPN IPv4 route to S and set the route target of the customer multicast route to the VRF import route carried by the VPN IPv4 route just found.

The process of originating customer multicast routes in an MBGP-based MVPN is shown in Figure 30 on page 361.

In the figure, an MVPN has three receiver sites (R1, R2, and R3) and one source site (S). The site routers are connected to four PE routers and PIM is running between the PE routers and the site routers. However, only BGP runs between the PE routers on the provider's network.

When router PE–1 receives a PIM join message for (S,G) from the site router R1, this means that site R1 has one or more receivers for a given source and multicast group (S,G) combination. In that case, router PE-1 constructs and originates a customer multicast route after doing three things:

1. Finding the unicast VPN IPv4 router to source S.
2. Extracting the route distinguisher and VRF route import form this route.
3. Putting the (S,G) information form the PIM join, the router distinguisher from the VPN IPv4 route and the route target from the VRF route import of the VPN IPv4 route into a MBGP update.

The update is distributed around the VPN through normal BGP mechanisms such as router reflectors.

Figure 30: Source and Receiver Sites in an MVPN

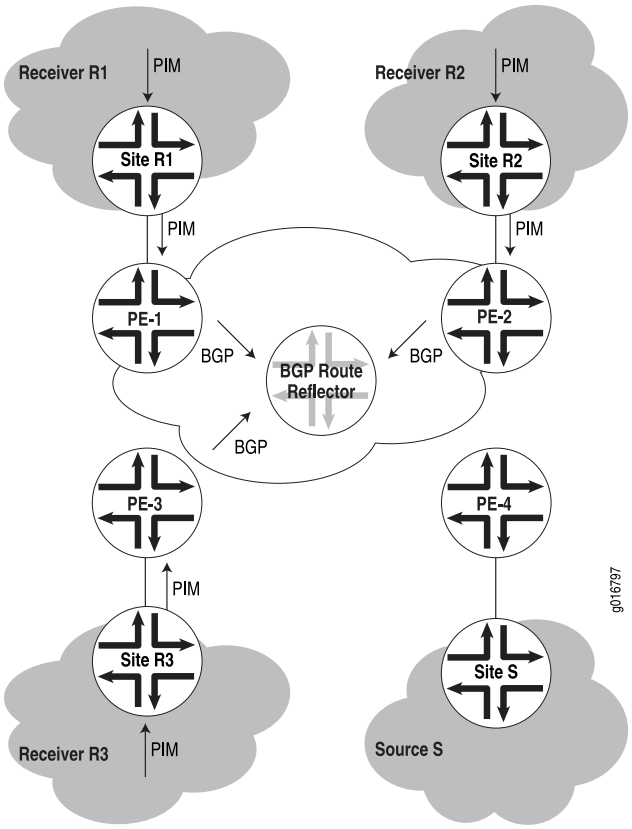

What happens when the source site S receives the MBGP information is shown in Figure 31 on page 362. In the figure, the customer multicast route information is distributed by the BGP route reflector as an MBGP update.

The provider router PE-4 will then:

1. Receive the customer multicast route originated by the PE routers and aggregated by the route reflector.
2. Accept the customer multicast route into the VRF for the correct MVPN (because the VRF route import matches the route target carried in the customer multicast route information).
3. Create the proper (S,G) state in the VRF and propagates the information to the customer routers of source site S using PIM.

Figure 31: Adding a Receiver to an MVPN Source Site Using MBGP

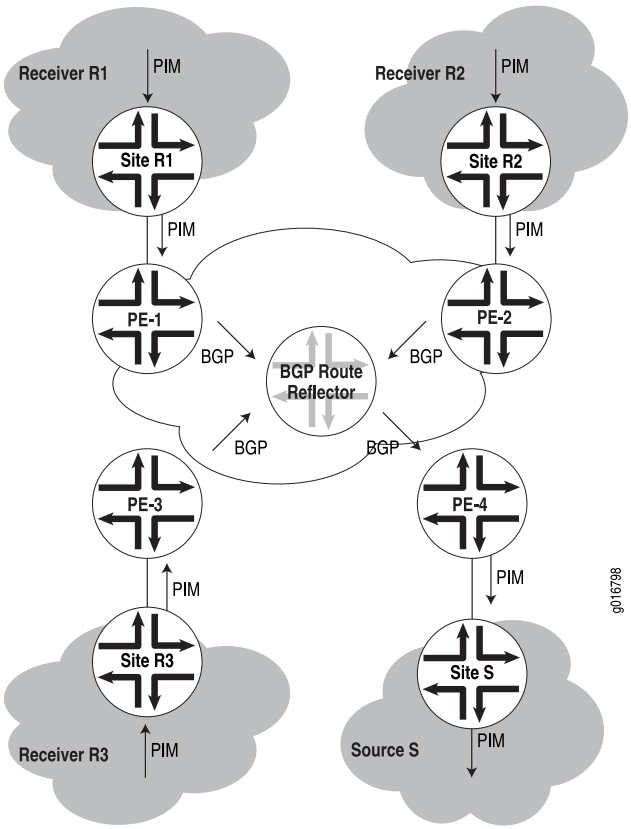

For more information about configuring multicast for Layer 3 VPNs using MBGP, see "Configuring Any-Source Multicast for Draft-Rosen VPNs" on page 415. For multicast Layer 3 VPN examples, see the *JUNOS Feature Guide*.

Tunnel Services PICs and Multicast

On Juniper Networks routers, data packets are encapsulated and de-encapsulated into tunnels by means of hardware and not the software running on the router's processor. The hardware used to create tunnel interfaces on M Series and T Series routers is a Tunnel Services PIC. If Juniper Networks M Series Multiservice Edge Routers and Juniper Networks T Series Core Routers are configured as rendezvous points or IP version 4 (IPv4) PIM sparse-mode DRs connected to a source, a Tunnel Services PIC is required. Juniper Networks MX Series Ethernet Services Routers do not require Tunnel Services PICs.

In PIM sparse mode, the source DR takes the initial multicast packets and encapsulates them in PIM register messages. It then unicasts them to the PIM sparse-mode RP router, where the PIM register message is de-encapsulated.

When a router is configured as a PIM sparse-mode RP router (by specifying an address using the address statement at the [edit protocols pim rp local] hierarchy level) and a Tunnel PIC is present on the router, a PIM register de-encapsulation interface, or pd

interface, is automatically created. The pd interface receives PIM register messages and de-encapsulates them by means of the hardware.

If PIM sparse mode is enabled on any router (potentially a PIM sparse-mode source DR) and a Tunnel Services PIC is present on the router, a PIM register encapsulation interface, or pe interface, is automatically created for each RP address that is used to encapsulate source data packets and send them to respective RP addresses on the PIM DR as well as the PIM RP. The pe interface receives PIM register messages and encapsulates them by means of the hardware.

NOTE: Do not confuse the configurable pe and pd hardware interfaces with the nonconfigurable pime and pimd software interfaces. Both pairs encapsulate and de-encapsulate multicast packets, and are created automatically; however, the pe and pd interfaces only appear if a Tunnel Services PIC is present. The pime and pimd interfaces are not useful in situations requiring the pe and pd interfaces.

If the source DR is the RP, then there is no need for PIM register messages and consequently no need for a Tunnel Services PIC to be present.

When PIM sparse mode is used with IP version 6 (IPv6), a Tunnel PIC is required on the RP, but not on the IPv6 PIM DR. The lack of a Tunnel PIC requirement on the IPv6 DR applies only to IPv6 PIM sparse mode and is not to be confused with IPv4 PIM sparse-mode requirements.

Table 7 on page 363 shows the complete matrix of IPv4 and IPv6 PIM Tunnel PIC requirements.

Table 7: Tunnel PIC Requirements for IPv4 and IPv6 Multicast

IP Version:	Tunnel PIC on RP	Tunnel PIC on DR
IPv4	Yes	Yes
IPv6	Yes	No

Filtering Multicast Messages

Multicast sources and routers generate a considerable number of control messages, especially when using PIM sparse mode. These messages form distribution trees, locate RPs and 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 two types of multicast filtering to control the sending and receiving of multicast control messages.

This section discusses three types of multicast filtering. The last two require configuration:

- Filtering MAC Addresses on page 364
- Filtering RP/DR Register Messages on page 364
- Filtering MSDP SA Messages on page 365

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 Open Shortest Path First (OSPF), Routing Information Protocol version 2 (RIPv2), or Network Time Protocol (NTP) are configured on an interface, each of these protocols individually requests the interface to program the MAC filter to pick up their respective multicast group alone. 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/DR Register Messages

You can filter PIM register messages sent from the DR or to the 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 knows about, or which sources a DR tells other RPs about, is desired. A high degree of control over PIM register messages is provided by RP/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 in combination.

If the action of the register filter policy is to discard the register message, the router needs to sends a register-stop message to the DR. These 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 are 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 will be 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. For more information on anycast RP, see “RP Mapping with Anycast RP” on page 357.

For more information on filtering RP/DR register messages, see “Configuring Register Message Filtering on a PIM RP or DR” on page 407 and “Example: Configuring Register Message Filters on PIM RPs and DRs” on page 428.

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 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 problems. See “Example: Configuring PIM BSR Filters” on page 427 for a sample filter configuration.

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.

For more information, see “Multicast Scoping Overview” on page 193 and “Example: Filtering Incoming PIM Join Messages” on page 427.

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.

Embedded RP for IPv6 Multicast

Global IPv6 multicast between routing domains has been possible only with SSM because there is no way to convey information about IPv6 multicast RPs between PIM sparse mode RPs. In IPv4 multicast networks this information is conveyed between PIM RPs using MSDP, but there is no IPv6 support in current MSDP standards. IPv6 uses the concept of an embedded RP to resolve this issue without requiring SSM. This feature embeds the RP address in an IPv6 multicast address.

All IPv6 multicast addresses begin with 8 1-bits (`1111 1111`) followed by a 4-bit flag field normally set to `0011`. The flag field is set to `0111` when embedded RP is used. Then the low-order bits of the normally reserved field in the IPv6 multicast address carry the 4-bit RP interface identifier (RIID).

When the IPv6 address of the RP is embedded in a unicast-prefix-based any-source multicast (ASM) address, all of the following conditions must be true:

- The address must be an IPv6 multicast address and have `0111` in the flags field (that is, the address is part of the prefix `FF70::/12`).
- The 8-bit prefix length (plen) field must not be all 0. An all 0 plen field implies that SSM is in use.
- The 8-bit prefix length field value must not be greater than 64, which is the length of the network prefix field in unicast-prefix-based ASM addresses.

The routing platform derives the value of the interdomain RP by copying the prefix length field number of bits from the 64-bit network prefix field in the received IPv6 multicast address to an empty 128-bit IPv6 address structure and copying the last bits from the 4-bit RIID. For example, if the prefix length field bits have the value 32, then the routing platform copies the first 32 bits of the IPv6 multicast address network prefix field to an all-0 IPv6 address and appends the last four bits determined by the RIID. See Figure 32 on page 367 for an illustration of this procedure.

Figure 32: Extracting the Embedded RP IPv6 Address

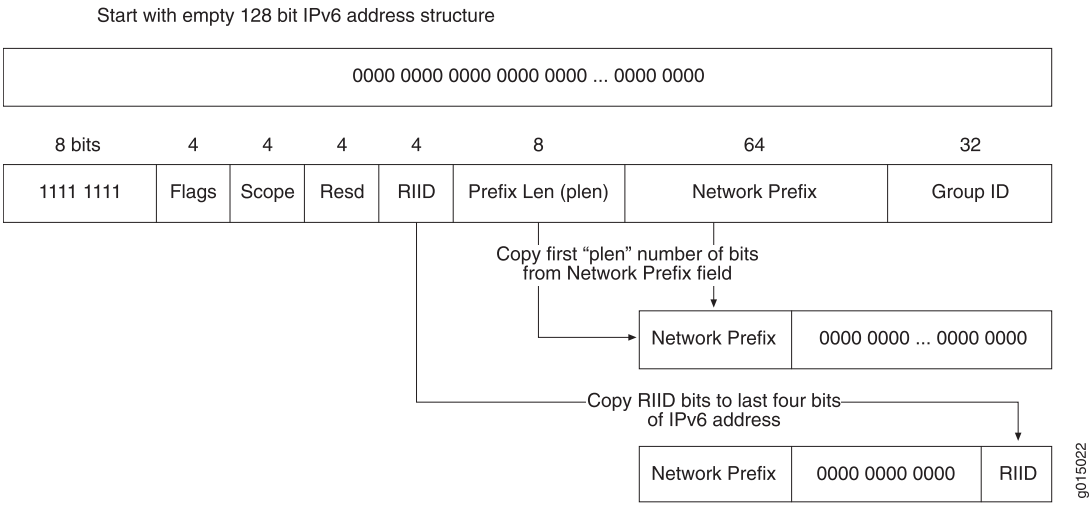

For example, the administrator of IPv6 network 2001:DB8::/32 sets up an RP for the 2001:DB8:BEEF:FEED::/96 subnet. In that case, the received embedded RP IPv6 ASM address has the form:

FF70:y40:2001:DB8:BEEF:FEED::/96

and the derived RP IPv6 address has the form:

2001:DB8:BEEF:FEED::y

where y is the RIID (y cannot be 0).

When configured, the routing platform checks for embedded RP information in every PIM join request received for IPv6. The use of embedded RP does not change the processing of IPv6 multicast and RPs in any way, except that the embedded RP address is used if available and selected for use. There is no need to specify the IPv6 address family for embedded RP configuration because the information can be used only if IPv6 multicast is properly configured on the routing platform.

The following receive events trigger extraction of an IPv6 embedded RP address on the routing platform:

- 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

The embedded RP node discovered through these events is added if it does not already exist on the routing platform. The routing platform chooses the embedded RP as the RP for a multicast group before choosing an RP learned through BSR or a statically configured RP. 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.

Chapter 39

PIM Mode-Independent Configuration Guidelines

This chapter discusses the following PIM features and configuration tasks:

- PIM Mode-Independent Configuration Overview on page 369
- PIM Configuration Statements on page 370
- Changing the PIM Version on page 372
- Configuring Interface Priority to Become the PIM Designated Router on page 372
- Configuring PIM Designated Router Election on Point-to-Point Links on page 373
- Modifying the PIM Hello Interval on page 373
- Configuring Interface-Level PIM Neighbor Policies on page 373
- Configuring a PIM RPF Routing Table on page 374
- Filtering Outgoing PIM Join Messages on page 376
- Filtering Incoming PIM Join Messages on page 377
- Preserving Multicast Performance by Disabling Response to the ping Utility on page 378
- Configuring the BFD Protocol for PIM on page 378
- Overview of BFD Authentication for PIM on page 379
- Configuring BFD Authentication for PIM on page 381
- Configuring Nonstop Active Routing with PIM on page 384
- Disabling PIM on page 385
- Configuring PIM Trace Options on page 388

PIM Mode-Independent Configuration Overview

This chapter discusses Protocol Independent Multicast (PIM) features that you can configuration on an interface regardless of its PIM mode (sparse, dense, or sparse-dense mode).

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. For

information about aggregate interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

Related Topics

- Configuring PIM Dense Mode Properties on page 391
- Configuring PIM Sparse-Dense Mode Properties on page 392
- Minimum PIM Sparse Mode Configuration on page 393
- Enabling PIM Sparse Mode on page 394
- Configuring PIM Sparse Mode Graceful Restart on page 395
- Configuring Local PIM RPs on page 396
- Configuring Static PIM RPs on page 399
- Configuring PIM Bootstrap Properties on page 400
- Configuring PIM Auto-RP on page 402
- Configuring Register Message Filtering on a PIM RP or DR on page 407
- Configuring PIM Join Load Balancing on page 408
- Modifying the IGMP Join State Timeout on page 412
- Configuring PIM Embedded RP for IPv6 on page 412
- Configuring the PIM Assert Timeout on page 413
- Configuring the PIM SPT Threshold Policy on page 413
- Configuring Any-Source Multicast for Draft-Rosen VPNs on page 415

PIM Configuration Statements

To configure PIM, include the pim statement:

```
pim {
    assert-timeout seconds;
    dense-groups {
        addresses;
    }
    disable;
    dr-election-on-p2p;
    export;
    graceful-restart {
        disable;
        restart-duration seconds;
    }
    import [ policy-names ];
    interface interface-name {
        import;
        hello-intervalseconds;
        mode (dense | sparse | sparse-dense);
        neighbor-policy [ policy-names ];
        priority number;
        version version;
    }
```

```
join-load-balance;
join-prune-timeout;
nonstop-routing {
   disable;
}
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-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;
   }
   local {
      family (inet | inet6) {
         address address;
         anycast-pim {
            rp-set {
               address address <forward-msdp-sa>;
            }
            local-address address;
         }
         disable;
         group-ranges {
            destination-ip-prefix</prefix-length>;
         }
         hold-time seconds;
         priority number;
      }
   }
   rp-register-policy [ policy-names ];
   static {
      address address {
         version version;
         group-ranges {
            destination-ip-prefix</prefix-length>;
         }
         spt-threshold {
            infinity [ policy-names ];
         }
         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-instance *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, PIM is disabled.

NOTE: You cannot configure PIM within a nonforwarding instance. If you try to do so, the router displays a commit check error and does not complete the configuration commit process.

Changing the PIM Version

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

By default, the JUNOS Software uses PIM version 2 (PIMv2). To configure PIM version 1 (PIMv1), include the version statement:

```
version 1;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

NOTE: 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.

Configuring Interface Priority to Become the PIM Designated Router

By default, a PIM interface has the lowest probability of being selected as the designated router (DR). To change this, include the priority statement:

```
priority number;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The default priority is 1. Use a larger number to increase the probability of the interface's being elected as the DR.

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

To comply with the latest PIM drafts, enable DR election on all PIM interfaces, including point-to-point (P2P) interfaces (DR election is enabled by default on all other interfaces). To enable DR election on P2P interfaces, include the dr-election-on-p2p statement:

```
dr-election-on-p2p;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Modifying the PIM Hello Interval

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

When a routing device is configured for PIM, it sends out 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 out another hello message, and the timer is reset. A routing device that gets 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 would wait for a response is 105 seconds.

If a PIM hello message is received that contains the holdtime option, the neighbor timeout is set to the holdtime sent in the message. If a PIM hello message is received that does not contain the holdtime option, the neighbor timeout is set to the default hello holdtime.

To modify how often the routing device sends hello messages out of an interface, include the hello-interval statement:

```
hello-interval seconds;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring Interface-Level PIM Neighbor Policies

To configure a policy to filter unwanted PIM neighbors, include the neighbor-policy statement:

```
neighbor-policy [ policy-names ];
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

The neighbor policy must be a properly structured policy statement that uses a prefix list containing the neighbor primary address (or any secondary IP addresses) in a prefix list and the reject option to reject the unwanted address. For example:

```
prefix-list nbrGroup1 {
   20.20.20.1/32;
}
policy-statement nbr-policy {
   from {
      prefix-list nbrGroup1;
   }
   then reject;
}
```

For more information about configuring prefix lists or policy statements, see the *JUNOS Policy Framework Configuration Guide*.

The PIM interface compares neighbor IP addresses with the IP addresses 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 with 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.

Configuring a PIM RPF Routing Table

By default, PIM uses inet.0 as its reverse-path-forwarding (RPF) routing table group. PIM uses an RPF routing table group to resolve its RPF neighbor for a particular multicast source address and to resolve the RPF neighbor for the RP address. PIM can optionally use inet.2 as its RPF routing table group. To do this, include the rib-groups statement at the [edit routing-options] hierarchy level. You must also include the export-rib statement because inet.0 cannot be the default table for exporting routes into another table with the import-rib statement. After configuring the routing table group, specify it as the value of the rib-group statement at the [edit protocols pim] hierarchy level:

```
[edit protocols pim]
rib-group group-name;
```

For more information on configuring RIB groups, see the *JUNOS Routing Protocols Configuration Guide*.

The following example defines the routing table group pim-rg and uses it to populate inet.2 for RPF checks:

```
[edit]
```

```
routing-options {
   rib-groups {
      pim-rg {
         export-rib inet.0;
         import-rib [ inet.0 inet.2 ];
      }
   }
}
protocols {
   pim {
      rib-group pim-rg;
   }
}
```

Specifying additional import routing table groups or an export routing table group in the routing table group has no effect on PIM operation. PIM uses the first routing table group specified as an import routing table group.

PIM uses a single routing table group as its RPF routing table group. This ensures that the route with the longest matching prefix is chosen as the RPF route.

You can configure OSPF to populate inet.2 with OSPF routes that have regular IP next hops. This allows RPF to work properly even when MPLS is configured for traffic engineering, or when OSPF is configured to use “shortcuts” for local traffic.

You can also configure IS-IS to populate inet.2 with IS-IS routes that have regular IP next hops. This allows RPF to work properly even when MPLS is configured for traffic engineering, or when IS-IS is configured to use “shortcuts” for local traffic.

For more information on RPF tables and the OSPF and IS-IS routing protocols, see the *JUNOS Routing Protocols Configuration Guide*.

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 a 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 will be 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 using 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 Topics Filtering Incoming PIM Join Messages on page 377

Filtering Incoming PIM Join Messages

While multicast scopes prevent 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 8 on page 377 for a list of match conditions.

Table 8: 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

To create a routing policy to reject a PIM join request for a source, include a policy name at the [edit policy-options policy-statement] or [edit logical-systems *logical-system-name* policy-options policy-statement] hierarchy level.

To apply one or more policies to routes being imported into the routing table from PIM, include the import statement:

```
import [ policy-names ];
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For a PIM join filter example, see "Example: Filtering Incoming PIM Join Messages" on page 427.

NOTE: Configuring multicast scoping on all routers filters the actual data and might be preferable to a PIM join filter solution. For more information about multicast scoping, see "Multicast Scoping Overview" on page 193.

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 configure, include the no-multicast-echo statement at the [edit system] hierarchy level:

```
system {
   no-multicast-echo;
}
```

For more information about this statement, see the *JUNOS System Basics Configuration Guide*.

Configuring the BFD Protocol for PIM

The bidirectional forwarding detection (BFD) protocol uses control packets and shorter detection time limits to detect failures more rapidly in a network. Working with a wide variety of network environments and topologies, BFD failure detection timers provide faster detection by using shorter time limits than the PIM hello hold time. These timers are also adaptive and you can adjust them to be more or less aggressive.

To enable failure detection, include the bfd-liveness-detection statement:

```
bfd-liveness-detection {
   detection-time {
      threshold milliseconds;
   }
   minimum-interval milliseconds;
   minimum-receive-interval milliseconds;
```

```
    multiplier number;
    version (0 | 1 | automatic);
}
```

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

To specify the minimum transmit and receive interval for failure detection, include the minimum-interval statement:

```
minimum-interval milliseconds;
```

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

To specify only the minimum receive interval for failure detection, include the minimum-receive-interval statement:

```
minimum-receive-interval milliseconds;
```

To specify the detection time multiplier for failure detection, include the multiplier statement:

```
multiplier number;
```

To specify the BFD version used for detection, include the version statement:

```
version (0 | 1 | automatic);
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

Overview of BFD Authentication for PIM

BFD enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when running BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with JUNOS Release 9.6, the JUNOS Software supports authentication for BFD sessions running over PIM. BFD authentication is only supported in the domestic image and is not available in the export image.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- BFD Authentication Algorithms on page 380
- Security Authentication Keychains on page 380
- Strict Versus Loose Authentication on page 381

BFD Authentication Algorithms

JUNOS Software supports the following algorithms for BFD authentication:

- simple-password—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords may be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- keyed-md5—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- meticulous-keyed-md5—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method may take additional time to authenticate the session.
- keyed-sha-1—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- meticulous-keyed-sha-1—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method may take additional time to authenticate the session.

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

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured

and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. 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.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Related Topics

- Configuring BFD Authentication for PIM on page 381
- bfd-liveness-detection statement
- authentication-key-chains statement in the *JUNOS System Basics Configuration Guide*
- show bfd session command in the *JUNOS Routing Protocols Command Reference*
- Configuring the BFD Protocol for PIM on page 378

Configuring BFD Authentication for PIM

Beginning with JUNOS Release 9.6, you can configure authentication for BFD sessions running over PIM. Routing instances are also supported. Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the PIM protocol.
2. Associate the authentication key chain 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 382
- Viewing Authentication Information for BFD Sessions on page 383

Configuring BFD Authentication Parameters

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]
user@host# set protocols pim interface if3-pim bfd-liveness-detection
   authentication algorithm keyed-sha-1
```

NOTE: Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms may 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. This should match the keychain name configured at the [edit security authentication key-chains] hierarchy level.

```
[edit]
user@host# set protocols pim interface if3-pim bfd-liveness-detection
   authentication keychain bfd-pim
```

NOTE: The algorithm and key chain 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 *key-chain-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, *YYYY-MM-DD.hh*:*mm*:*ss*.

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

```
user@host> set protocols pim interface if3-pim bfd-liveness-detection
  authentication loose-check
```

5. (Optional) View your configuration 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.

NOTE: BFD authentication is only supported in the domestic image and is not available in the export image.

Viewing Authentication Information for BFD Sessions

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

The following example shows BFD authentication configured for the if3–pim BGP group. It specifies the keyed SHA-1 authentication algorithm and a key-chain name of bfd-pim. The authentication keychain is configured with two keys. Key 1 contains the secret data "9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm" and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data "9a5jiKW9I.reP38ny.TszF2/9" and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols pim]
interface if3–pim {
   bfd-liveness-detection {
      authentication {
         algorithm keyed-sha-1;
         key-chain bfd-pim;
      }
   }
}
[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 "$9$a5jiKW9I.reP38ny.TszF2/9";
         start-time "2009-6-1.15:29:20 -0700";
      }
   }
}
```

If you commit these updates to your configuration, you would see output similar to the following. In the output for the show bfd sessions detail command, Authenticate is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the show bfd sessions 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 sessions detail

```
user@host# show bfd session detail

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

show bfd sessions extensive

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

Related Topics

- Overview of BFD Authentication for PIM on page 379
- bfd-liveness-detection statement
- authentication-key-chains statement in the *JUNOS System Basics Configuration Guide*
- show bfd session command in the *JUNOS Routing Protocols Command Reference*
- Configuring the BFD Protocol for PIM on page 378

Configuring Nonstop Active Routing with PIM

Nonstop active routing configurations include two Routing Engines that share information so that routing is not interrupted during Routing Engine failover. When nonstop active routing is configured on a dual Routing Engine platform, the PIM control state is replicated on both Routing Engines.

This PIM state information includes:

- Neighbor relationships
- Join and prune information

- RP-set information
- Synchronization between routes and next hops and the forwarding state between the two Routing Engines

The PIM control state is maintained on the backup Routing Engine by the replication of state information from the master to the backup Routing Engine and having the backup Routing Engine react to route installation and modification in the [instance].inet.1 routing table on the master Routing Engine. The backup Routing Engine does not send or receive any PIM protocol packets directly. In addition, the backup Routing Engine uses the dynamic interfaces created by the master Routing Engine. These dynamic interfaces include PIM encapsulation, de-encapsulation, and multicast tunnel interfaces.

NOTE: The clear pim join, clear pim register, and clear pim statistics operational mode commands are not supported on the backup Routing Engine when nonstop active routing is enabled.

To enable nonstop active routing for PIM (in addition to the PIM configuration on the master Routing Engine), you must include the following statements at the [edit] hierarchy level:

- chassis redundancy graceful-switchover
- routing-options nonstop-routing
- system commit synchronize

For more information about PIM and nonstop active routing, see "IGMP and Nonstop Active Routing" on page 73 and the *JUNOS High Availability Configuration Guide*.

Disabling PIM

By default, when configured, the PIM protocol is enabled on all interfaces for all families. If desired, you may 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 386
- Disabling PIM on an Interface on page 386
- Disabling PIM for a Family on page 387
- Disabling PIM for a Rendezvous Point on page 388

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:

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

To disable the PIM protocol:

1. Include the disable statement:

```
user@host# set protocols pim disable
```

2. You can verify your configuration settings prior to committing them using the show protocols 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:

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

To disable PIM for an interface:

1. Include the disable statement:

```
user@host# set protocols pim interface fe-0/1/0 disable
```

2. You can verify your configuration settings prior to committing them 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:

```
[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. You can verify your configuration settings prior to committing them 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 rendezvous point (RP) on a per-family basis. This is accomplished at the [edit protocols pim rp local family] hierarchy:

```
[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. You can verify your configuration settings prior to committing them using the show protocols pim command:

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

Configuring PIM Trace Options

To trace PIM protocol traffic, you can specify options in the global traceoptions statement at the [edit routing-options] or [edit logical-systems *logical-system-name* routing-options] hierarchy level. Options applied at the routing options level trace all packets, and options applied at the protocol level trace only IGMP traffic.

```
traceoptions {
   file filename <files number> <size size> <world-readable | no-world-readable>;
   flag flag <flag-modifier> <disable>;
}
```

You can configure PIM-specific options by including the traceoptions statement at the PIM hierarchy level. For a list of the hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

You can specify the following PIM-specific options in the traceoptions statement:

- 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.
- 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.
- cache—Trace the packets in the PIM routing cache.
- graft—Trace graft and graft acknowledgment messages.
- hello—Trace hello packets, which are sent so that neighboring routers can discover each other.
- join—Trace join messages, which are sent to join a branch onto the multicast distribution tree.
- packets—Trace all PIM 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.
- rp—Trace candidate RP advertisements.

For general information about tracing, see the *JUNOS System Basics Configuration Guide*. For a PIM tracing example, see “Example: Tracing PIM Protocol Traffic” on page 430.

Chapter 40

PIM Dense and Sparse-Dense Mode Configuration Guidelines

This chapter explains how to configure Protocol Independent Multicast (PIM) dense mode and sparse-dense mode properties:

- Configuring PIM Dense Mode Properties on page 391
- Configuring PIM Sparse-Dense Mode Properties on page 392

Configuring PIM Dense Mode Properties

To configure the router properties for PIM dense mode, enable the minimum PIM dense mode configuration. For information about operating interfaces in PIM dense mode, see "PIM Modes of Operation" on page 28.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default. To enable PIM dense mode on the router, include the pim statement:

```
pim {
    rib-group group-name;
    interface interface-name {
        mode dense;
    }
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To specify that PIM dense mode use inet.2 as its RPF routing table instead of inet.0, include the rib-group statement. For more information about configuring RPF routing tables, see "Configuring a PIM RPF Routing Table" on page 374.

You can specify the interfaces on which to enable PIM. Specify the full name, including the physical and logical address components. For details about specifying interfaces, see the *JUNOS Network Interfaces Configuration Guide*. 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.

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.

Configuring PIM Sparse-Dense Mode Properties

To configure PIM to operate in sparse-dense mode on an interface, include the mode sparse-dense statement. Include the dense-groups statement to specify which groups are operating in dense mode:

```
dense-groups {
    addresses;
}
interface interface-name {
    mode sparse-dense;
}
```

For a list of the hierarchy levels at which you can include these statements, see the statement summary section for this statement.

You can configure graceful restart with PIM sparse-dense mode, but only sparse multicast groups benefit from graceful restart. For more information about graceful restart for PIM sparse mode, see “Configuring PIM Sparse Mode Graceful Restart” on page 395.

For an example of a sparse-dense mode configuration, see “Example: Configuring PIM Sparse-Dense Mode” on page 422.

Chapter 41

PIM Sparse Mode Configuration Guidelines

This chapter explains how to configure Protocol Independent Multicast (PIM) sparse mode properties. For information about operating interfaces in PIM sparse mode, see "PIM Modes of Operation" on page 28.

- Minimum PIM Sparse Mode Configuration on page 393
- Enabling PIM Sparse Mode on page 394
- Configuring PIM Sparse Mode Graceful Restart on page 395
- Configuring Local PIM RPs on page 396
- Configuring Static PIM RPs on page 399
- Configuring PIM Bootstrap Properties on page 400
- Configuring PIM Auto-RP on page 402
- Configuring Register Message Filtering on a PIM RP or DR on page 407
- Configuring PIM Join Load Balancing on page 408
- Modifying the IGMP Join State Timeout on page 412
- Configuring PIM Embedded RP for IPv6 on page 412
- Configuring the PIM Assert Timeout on page 413
- Configuring the PIM SPT Threshold Policy on page 413
- Configuring Any-Source Multicast for Draft-Rosen VPNs on page 415

Minimum PIM Sparse Mode Configuration

Each any-source multicast (ASM) group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The RP is the root of this shared tree.

To configure this router's properties as the candidate rendezvous point (RP), include the rp statement:

```
rp {
   local {
      family (inet | inet6) {
         disable;
         address address;
```

```
            group-ranges {
                destination-ip-prefix</prefix-length>;
            }
            hold-time seconds;
            priority number;
        }
    }
    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-export [ policy-names ];
    bootstrap-import [ policy-names ];
    bootstrap-priority number;
    dr-register-policy [ policy-name ];
    embedded-rp {
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
        maximum-rps limit;
    }
    rp-register-policy [ policy-name ];
    static {
        address address {
            version version;
            group-ranges {
                destination-ip-prefix</prefix-length>;
            }
        }
    }
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Enabling PIM Sparse Mode

You can configure PIM interfaces to operate in sparse, dense, or sparse-dense mode. Sparse mode is the default. There is no need to explicitly configure sparse mode on a PIM interface, but this is often done for clarity or when you configure a change from dense to sparse mode.

To explicitly configure PIM to operate in sparse mode on an interface, include the mode sparse statement:

```
mode sparse;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring PIM Sparse Mode Graceful Restart

You can configure PIM sparse mode to continue to forward existing multicast packet streams during a routing process failure and restart. Only PIM sparse mode can be configured this way. The routing platform does not forward multicast packets for protocols other than PIM during graceful restart, because all other multicast protocols must restart after a routing process failure.

NOTE: If you configure PIM sparse-dense mode, only sparse multicast groups benefit from graceful restart.

The routing platform does not forward new streams until after the restart is complete. After restart, the routing platform refreshes the forwarding state with any updates that were received from neighbors during the restart period. For example, the routing platform relearns the join and prune states of neighbors during the restart, but it does not apply the changes to the forwarding table until after the restart.

When PIM sparse mode is enabled, the routing platform generates a unique 32-bit random number called a generation identifier. Generation identifiers are included by default in PIM hello messages, as specified in the Internet draft Internet draft draft-ietf-pim-sm-v2-new-10.txt. When a routing platform receives PIM hello messages containing generation identifiers on a point-to-point interface, the JUNOS Software activates an algorithm that optimizes graceful restart.

Before PIM sparse mode graceful restart occurs, each routing platform creates a generation identifier and sends it to its multicast neighbors. If a routing platform with PIM sparse mode restarts, it creates a new generation identifier and sends it to neighbors. When a neighbor receives the new identifier, it resends multicast updates to the restarting router to allow it to exit graceful restart efficiently. The restart phase is complete when the restart interval timer expires. On platforms that support PIM sparse mode graceful restart, the restart can be completed within 30 through 300 seconds. The default restart duration is 60 seconds.

NOTE: Multicast forwarding can be interrupted in two ways. First, if the underlying routing protocol is unstable, multicast RPF checks can fail and cause an interruption. Second, because the forwarding table is not updated during the graceful restart period, new multicast streams are not forwarded until graceful restart is complete.

To configure graceful restart for PIM sparse mode, include the graceful-restart statement:

```
graceful-restart {
    disable;
    restart-duration seconds;
}
```

For a list of the hierarchy levels at which you can configure the graceful-restart statement, see the statement summary section for this statement.

To disable graceful restart for PIM, include the disable statement:

```
disable;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

By default, the router allows 60 seconds for the restart duration. The range is from 30 through 300 seconds. After this restart time, the Routing Engine resumes normal multicast operation. To configure the restart duration, include the restart-duration statement:

```
restart-duration seconds;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For more information about graceful restart for PIM, see "Multicast Redundancy" on page 20. For more information about graceful restart and other routing protocols, see the *JUNOS High Availability Configuration Guide*, *JUNOS Routing Protocols Configuration Guide* and the *JUNOS Feature Guide*.

Configuring Local PIM RPs

Local RP configuration makes the routing device a statically defined RP. To configure the routing device's RP properties, include the local statement:

```
local {
   family (inet | inet6) {
      disable;
      address address;
      group-ranges {
      destination-ip-prefix</prefix-length>;
      }
      hold-time seconds;
      priority number;
   }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols pim rp]
- [edit routing-instances *routing-instance-name* protocols pim rp]

For an overview of logical systems and a detailed example of logical system configuration, see the logical systems chapter of the *JUNOS Feature Guide*.

For information about the RP configuration statements, see the following sections:

- Configuring the IP Protocol Family on page 397
- Configuring the Local RP Address on page 398
- Configuring the Routing Device's RP Priority on page 398
- Configuring the Groups for Which the Routing Device Is the RP on page 398
- Modifying the Local RP Hold Time on page 399

Configuring the IP Protocol Family

PIM supports both IP version 4 (IPv4) and IP version 6 (IPv6) addressing.

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.

For correct operation of PIM sparse mode, the RP address must be known to a routing device. The JUNOS IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6. You configure the static IPv6 RP address in the same way as IPv4 addresses, by including the address statement. However, on a routing device that is itself the RP, include the address statement at the [edit protocols pim rp local family inet6] or [edit routing-instances *routing-instance-name* protocols pim rp local family inet6] hierarchy level.

For an overview of logical systems and a detailed example of logical system configuration, see the logical systems chapter of the *JUNOS Feature Guide*.

The Multicast Listener Discovery (MLD) protocol is automatically enabled on any broadcast interfaces on which you configure PIM and family inet6. For an overview of MLD, see "MLD Overview" on page 95.

To specify whether IPv4 or IPv6 local RP properties apply to the configuration values, include the family statement:

```
family (inet | inet6);
```

You can include this statement at the following hierarchy levels:

- [edit protocols pim rp local]
- [edit routing-instances *routing-instance-name* protocols pim rp local]

Configuring the Local RP Address

To specify the local RP address, include the address statement:

```
address address;
```

You can include this statement at the following hierarchy levels:

- [edit protocols pim rp local family]
- [edit routing-instances *routing-instance-name* protocols pim rp local family]

Configuring the Routing Device's RP Priority

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 well-known 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. To modify the routing device's priority, include the priority statement:

```
priority number;
```

You can include this statement at the following hierarchy levels:

- [edit protocols pim rp local family]
- [edit routing-instances *routing-instance-name* protocols pim rp local family]

The priority can be a number from 0 through 255.

Configuring 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 groups (224.0.0.0/4). To limit the groups for which this routing device can be the RP, include the group-ranges statement:

```
group-ranges {
    destination-ip-prefix</prefix-length>;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols pim rp local family]
- [edit routing-instances *routing-instance-name* protocols pim rp local family]

Modifying 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.

To modify the hold-time value for the local RP, include the hold-time statement:

```
hold-time seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols pim rp local family]
- [edit routing-instances *routing-instance-name* protocols pim rp local family]

Configuring Static PIM RPs

Static RP configuration directs the routing device to another statically defined RP. To configure static RPs, include the static statement:

```
static {
   address address {
      version version;
      group-ranges {
         destination-ip-prefix</prefix-length>;
      }
   }
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

NOTE: You can configure a static RP in a logical system only if the logical system is not directly connected to a source.

To configure other static RPs, include one or more address statements. The default multicast address group range is 224.0.0.0/4.

For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.

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

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 (at the [edit pim rp static address *address*] hierarchy level). However, PIM version 2 is the default for interface mode (at the [edit pim interface *interface-name*] hierarchy level). Explicitly configured versions override the defaults.

Configuring PIM Bootstrap Properties

Bootstrap routers are supported in IPv4 and IPv6. For legacy configuration purposes, the first two of the following sections describe configuration of IPv4 bootstrap routers. However, a different configuration hierarchy can be used for both IPv4 and IPv6, as described in the third and fourth sections—this combined configuration method is recommended.

- Configuring the Routing Device's IPv4 Bootstrap Router Priority on page 400
- Filtering PIM IPv4 Bootstrap Messages on page 400
- Configuring the Routing Device's Bootstrap Router Priority on page 401
- Filtering PIM Bootstrap Messages on page 402

Configuring the Routing Device's IPv4 Bootstrap Router Priority

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 are configured to operate within a common boundary. The domain's bootstrap router originates 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.

By default, the routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap router. To modify this priority, include the bootstrap-priority statement. 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 simply assigns a bootstrap priority value to a routing device.

```
bootstrap-priority number;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Filtering PIM IPv4 Bootstrap Messages

You can create import and export policies to control the flow of IPv4 bootstrap messages to and from the RP, and apply them to PIM. To apply one or more import policies to IPv4 bootstrap messages imported into the RP, include the bootstrap-import statement:

```
bootstrap-import [ policy-names ];
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To apply one or more export policies to IPv4 bootstrap messages exported from the RP, include the bootstrap-export statement:

bootstrap-export [*policy-names*];

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

NOTE: The bootstrap-priority, bootstrap-import, and bootstrap-export statements support IPv4 only. A priority of 0 disables the function for IPv4 and does not cause the routing device to send BSR packets with a 0 in the priority field.

Configuring the Routing Device's Bootstrap Router Priority

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 are configured to operate within a common boundary. The domain's bootstrap router originates 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.

The bootstrap configuration hierarchy supports both IPv4 and IPv6 multicasting. It can be combined with the bootstrap statements supported in IPv4 only, as long as the added statements are used for IPv6 only, but this is not recommended. There is a change in the meaning of the bootstrap priority when the value is set to 0.

In the IPv4 configuration hierarchy, specifying the value 0 (zero) for the bootstrap-priority statement disables the function for IPv4 and does not cause the routing device to send BSR packets with a 0 in the priority field (for more information, see "Configuring the Routing Device's IPv4 Bootstrap Router Priority" on page 400). In the combined IPv4 and IPv6 configuration hierarchy, specifying the value 0 for the priority statement 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 hierarchy, delete the configuration statements.

A commit error occurs if the same IPv4 bootstrap statements are included in both the IPv4-only and the IPv4-and-IPv6 sections of the hierarchy. The error message is "duplicate IPv4 bootstrap configuration."

We recommend that legacy IPv4-only configurations be converted to the IPv4 and IPv6 configuration hierarchy.

To modify the bootstrap priority for IPv4 or IPv6, include the priority statement for the appropriate address family: inet for IPv4 and inet6 for IPv6. 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.

priority *number*;

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The following example sets bootstrap priority to 1 for both IPv4 and IPv6 multicasts:

```
[edit protocols pim rp]
bootstrap {
   family inet {
      priority 1;
   }
   family inet6 {
      priority 1;
   }
}
```

An error results when the configuration at this hierarchy level includes both these statements and the bootstrap-priority statement.

Filtering PIM Bootstrap Messages

You can create import and export policies to control the flow of bootstrap messages to and from the RP, and apply them to PIM. To apply one or more import policies to bootstrap messages imported into the RP, include the import statement:

import [*policy-names*];

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To apply one or more export policies to bootstrap messages exported from the RP, include the export statement:

export [*policy-names*];

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For an example, see "Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain" on page 431.

NOTE: The bootstrap statement stanza supports both IPv4 and IPv6. A 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, delete the configuration statements.

Configuring PIM Auto-RP

You can configure a mode-dynamic way of assigning RPs in a multicast network by means of auto-RP. When you configure auto-RP for a routing device, the routing device learns the address of the RP in the network automatically. Auto-RP operates in PIM version 1 and version 2.

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.

To configure auto-RP properties, see the following sections:

- Configuring Auto-RP Announcement, Mapping, and Discovery on page 403
- Configuring Auto-RP Mapping Agent Election on page 406

Configuring Auto-RP Announcement, Mapping, and Discovery

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: you can configure multiple routing devices as RP candidates. If the elected RP stops operating, one of the other preconfigured routing devices takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

If PIM is operating in sparse or sparse-dense mode, configure how the routing device operates in auto-RP by specifying the following auto-RP options:

- Add the discovery option to enable the routing device to receive and process discovery messages from the mapping agent. This is the most basic auto-RP option.
- Add the announce option to enable the routing device to send announce messages in the network, advertising itself as a candidate RP. Routing devices configured with this option must also be configured as RPs, or announce messages are not sent.
- Add the mapping option to enable the routing device to perform the mapping function. If the routing device 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).

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. The router ID is used as the address for auto-RP updates. 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.

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 403 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.

Table 9: Local RP and Auto-RP Message Types *(continued)*

Auto-RP Message Type	Local RP?	Routing Device Behavior
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.

To configure auto-RP at the main hierarchy level, follow these steps:

1. Include the mode statement, and specify the option sparse-dense on all interfaces at the [edit protocols pim] hierarchy level:

   ```
   [edit protocols pim]
   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 two multicast dense mode groups (224.0.1.39 and 224.0.1.40) using the dense-groups statement at the [edit protocols pim] hierarchy level:

   ```
   [edit protocols pim]
   dense-groups {
      224.0.1.39/32;
      224.0.1.40/32;
   }
   ```

   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.

3. Include the auto-rp statement at the [edit protocols pim rp] hierarchy level to configure auto-RP on each routing device in the group. There are four possible categories for each routing device.

   - If the routing device is not a local RP and listens only for auto-RP mapping messages, include the auto-rp discovery statement at the [edit protocols pim rp] hierarchy level:

```
[edit protocols pim rp]
auto-rp discovery;
```

- If the routing device is a local RP, sends auto-RP announcements, and listens for auto-RP mapping messages, configure the routing device as a local RP and include the **auto-rp announce** statement to the routing device RP configuration at the **[edit protocols pim rp]** hierarchy level:

```
[edit protocols pim rp]
local {
   address 10.0.1.1;
}
auto-rp announce;
```

- If the routing device performs only the mapping function to listen for auto-RP announcements, performs the auto-RP-to-group mapping, and sends auto-RP mapping messages, include the **auto-rp mapping** statement at the **[edit protocols pim rp]** hierarchy level. When multiple candidate RP routing devices announce their capabilities to support multicast groups, there must be a single routing device in the network to act as mapping agent. The mapping agent sends out discovery messages to the network, informing all routing devices in a multicast group of the RP to use:

```
[edit protocols pim rp]
auto-rp mapping;
```

- If the routing device combines the local RP function to send announcements and also perform the mapping function, configure the routing device as a local RP and include the **auto-rp mapping** statement at the **[edit protocols pim rp]** hierarchy level:

```
[edit protocols pim rp]
local {
   address 10.0.1.1;
}
auto-rp mapping;
```

All routing devices must also have a routable IP address on the **lo0** interface:

```
[edit]
interfaces lo0 {
   unit 0 {
      family inet {
         address 127.0.0.1; # this address cannot be used by auto-rp
         address 192.168.27.1 { # this example uses a private IP address
            preferred;
         }
      }
   }
}
```

You can include these statements at the following hierarchy levels (auto-RP announce is not supported in logical systems):

- [edit protocols pim]
- [edit routing-instances *routing-instance-name* protocols pim]
- [edit logical-systems *logical-system-name* protocols pim] (all statements except auto-rp announce)
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim] (all statements except auto-rp announce)

To verify auto-RP information, issue the show pim rps command :

```
user@host> show pim rps
RP address     Type    Holdtime    Timeout  Active groups   Group prefixes
192.168.5.1    auto-rp     150       123                1   224.0.0.0/4
```

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.

Configuring Auto-RP Mapping Agent Election

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. This process is called *mapping agent election*. 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.

For compatibility, you can suppress mapping messages with 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.

The default auto-RP operation is to perform mapping agent election. To explicitly enable mapping agent election, include the mapping-agent-election statement for an auto-RP mapping agent:

```
[edit protocols pim rp]
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 explicitly disable mapping agent election for compatibility with other vendors' equipment, include the no-mapping-agent-election statement for an auto-RP mapping agent:

```
[edit protocols pim rp]
auto-rp {
   mapping;
   no-mapping-agent-election;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols pim rp]
- [edit routing-instances *routing-instance-name* protocols pim rp]
- [edit logical-systems *logical-system-name* protocols pim rp]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp]

Configuring Register Message Filtering on a PIM RP or DR

You configure RP/DR register message filtering to control the number and location of multicast sources that an RP knows. 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.

To filter incoming register messages at the RP, include the rp-register-policy statement for the RP:

```
rp {
   rp-register-policy [ rp-register-policies ];
   local {
      address address;
```

```
    }
}
```

To filter outgoing register messages at the DR, configure the **dr-register-policy** statement for the DR:

```
rp {
    dr-register-policy [ policy-names ];
    static {
        address address;
    }
}
```

You can include these statements at the following hierarchy levels:

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

If you delete a group and source address from a filter policy on an RP router, the RP will register the group and source only when the DR sends a null register message.

You can configure more than one policy for each statement. If a referenced policy does not exist, the configuration commit checkout will fail. For examples of both types of register filters, see "Example: Configuring Register Message Filters on PIM RPs and DRs" on page 428.

For more information on RP/DR register message filtering, see "Filtering RP/DR Register Messages" on page 364.

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 and add to the multicast data tree. This interface is also used for all downstream traffic, so even though there are alternatives 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 towards a source. PIM join load balancing is only supported for PIM sparse mode configurations.

NOTE: 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).

NOTE: PIM join load balancing is not supported and should not be configured when protocol-independent load balancing for Layer 3 VPNs is configured using the multipath statement.

To configure join load balancing for PIM sparse mode, include the join-load-balance statement:

```
join-load-balance;
```

You can include this statement at the following hierarchy levels:

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

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 and not on a per-source basis. Therefore, there is no guarantee that joins for a particular source will be load-balanced, but the joins for all sources and all groups known to the routing device will be 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.

This example configures join load balancing for PIM sparse mode:

```
[edit protocols pim]
rp {
   static {
      address 10.10.10.1;
   }
}
interface all {
```

```
    mode sparse;
    version 2;
  }
  join-load-balance;
```

You can find out 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.

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 join extensive command now have nonzero join counts. If the counts differ by more than one and were zero when load balancing commenced, there is an error (joins prior to load balancing are not redistributed). 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.

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.

Modifying the IGMP 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 message. Because the messages are called join-prune messages, the name used is the join-prune-timeout statement.

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

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

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

Related Topics

- join-prune-timeout

Configuring PIM Embedded RP for IPv6

You configure embedded RP to allow multidomain IPv6 multicast networks to find RPs in other routing domains. Embedded RP embeds an RP address inside PIM join messages and other types of messages sent between routing domains.

Embedded RP is disabled by default. To configure embedded RP for IPv6 PIM sparse mode, include the embedded-rp statement:

```
embedded-rp {
   group-ranges {
      destination-ip-prefix</prefix-length>;
   }
   maximum-rps limit;
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

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.

The group-ranges statement determines which multicast addresses or prefixes can embed RP address information. If messages within a group range contain embedded RP information and the group range is not configured, the embedded RP in that group range is ignored. Any valid unicast-prefix-based ASM address can be used as a group range. The default group range is FF70::/12 to FFF0::/12. Messages with embedded RP information that do not match any configured group ranges are treated as normal multicast addresses.

If the derived RP address is not a valid IPv6 unicast address, it is treated as any other multicast group address and not used for RP information. Verification fails if the extracted RP address is a local interface, unless the routing device is configured as an RP and the extracted RP address matches the configured RP address. Then the local RP decides whether it is configured to act as an RP for the embedded RP multicast address.

When you configure embedded RP for IPv6, embedded RPs are preferred to RPs discovered by IPv6 any other way. You configure embedded RP independent of any other IPv6 multicast properties. This feature is applied only when IPv6 multicast is properly configured.

For more information about the use of embedded RP, see "Embedded RP for IPv6 Multicast" on page 366.

Configuring the PIM Assert Timeout

You configure the assert timeout to determine how often multicast routing devices running PIM sparse mode enter a PIM assert message cycle. 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 routers determine which routing device forwards the traffic and prunes the RPT for this group.

Assert messages are useful for LANs that connect multiple routing devices and no hosts. For more information about network configurations using assert timeouts, see "SPT Cutover" on page 350 and "SPT Cutover Control" on page 354.

To configure the assert timeout for PIM sparse mode, include the assert-timeout statement:

```
assert-timeout seconds;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

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

Configuring the PIM SPT Threshold Policy

Multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through a rendezvous-point tree (RPT) rooted at the RP or through a shortest-path tree (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. For more information about these SPT cutover cases, see "SPT Cutover Control" on page 354.

You configure an SPT threshold policy on the last-hop routing device 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 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.

To configure the SPT threshold and policy for PIM sparse mode, include the spt-threshold statement:

```
spt-threshold {
   infinity [ policy-names ];
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

NOTE: If you want the PE router to stay on the RPT for control traffic, include the spt-threshold statement under the main PIM instance.

The infinity 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.

Several points are important when configuring the SPT threshold policy, as discussed in the following sections:

- SPT Threshold Policy Configuration Changes on page 414
- Examples of SPT Threshold Policy Configuration on page 414

SPT Threshold Policy Configuration Changes

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.

NOTE: 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.

Examples of SPT Threshold Policy Configuration

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.

```
protocols {
   pim {
      ...
      spt-threshold {
         infinity spt-threshold-infinity-policy;
      }
      ...
   }
   ...
   policy-options {
      policy-statement spt-threshold-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;
               }
            }
         }
      }
   }
}
```

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 will never transition 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 will transition to a direct SPT to the source.

Configuring Any-Source Multicast for Draft-Rosen VPNs

If the service provider supports PIM, you can configure multicast for a Layer 3 virtual private network (VPN) using PIM version 2 as the multicast protocol. The JUNOS Software complies with RFC 4364 and Internet draft draft-rosen-vpn-mcast-07.txt, *BGP/MPLS VPNs* (known as draft-rosen) and *Multicast in MPLS/BGP VPNs*.

This section describes configuration of an any-source multicast VPN (MVPN) using dual PIM configuration with a customer RP and provider RP and mapping the multicast routes from customer to provider (known as draft-rosen). You can also configure a MVPN using Multiprotocol BGP (MBGP), or next generation MVPN, which is described in the next section. In this section, the term *Layer 3 VPNs* is used to refer to draft-rosen MVPNs.

For multicast to work on draft-rosen Layer 3 VPNs, each of the following routers must have a Tunnel Services PIC, which is hardware used to encapsulate and de-encapsulate data packets into tunnels:

- Each provider edge (PE) router
- Any provider (P) router acting as the RP

- Any customer edge (CE) router that is acting as a source's DR or as an RP. A receiver's designated router does not need a Tunnel Services PIC.

When you complete the configuration, two multicast tunnel interfaces are configured automatically. You do not need to configure the tunnel interfaces. The interface mt-[*xxxxx*], used for encapsulation, is in the range from 32,768 through 49,151. The interface mt-[*yyyyy*], used for de-encapsulation, is in the range from 49,152 through 65,535. For each VPN, the PE routers build a multicast distribution tree within the service provider core network. After the tree is created, each PE router encapsulates all multicast traffic (data and control messages) from the attached VPN and sends the encapsulated traffic to the VPN group address. Because all the PE routers are members of the outgoing interface list in the multicast distribution tree for the VPN group address, they all receive the encapsulated traffic. When the PE routers receive the encapsulated traffic, they de-encapsulate the messages and send the data and control messages to the CE routers.

NOTE: It is possible for the PE router to be configured as the VPN customer RP (C-RP) router. The PE router can also act as the DR. This type of PE configuration can simplify configuration of customer DRs and VPN C-RPs for multicast VPNs. However, the BSR and auto-RP features are not supported. This section does not discuss the use of the PE as the VPN C-RP.

The following sections describe how to configure multicast for Layer 3 VPNs:

- Configuring the VPN on page 416
- Configuring PIM Connectivity Between the Provider and PE Routers on page 417
- Configuring Multicast Connectivity on the CE Routers on page 417
- Configuring Multicast Connectivity for the VPN on the PE Router on page 418
- Configuring the Routing Group on page 418

Configuring the VPN

You must first configure the VPN. Figure 33 on page 416 shows a configuration for VPN-A, used as an example later in this section. For more information about configuring VPNs, see the *JUNOS VPNs Configuration Guide*.

Figure 33: Configuring the VPN

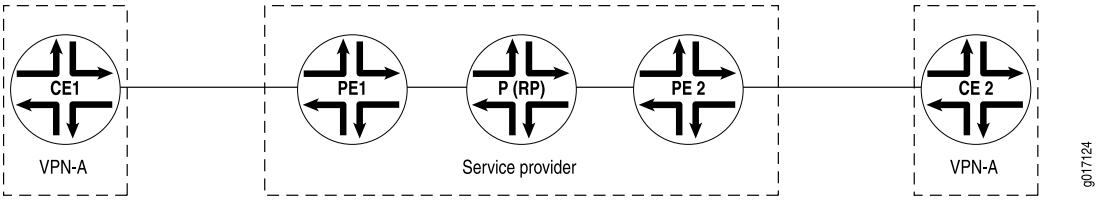

Configuring PIM Connectivity Between the Provider and PE Routers

To configure PIM on the main routing instance for all provider and PE routers, include statements at the [edit protocols pim] hierarchy level:

1. Configure the interfaces between each provider router and the PE routers by including the interface statement at the [edit protocols pim] hierarchy level. On all PE routers, enable PIM version 2 and sparse mode on interface lo0 of the PE routers, either by configuring that specific interface or by including the statement set version 2 mode sparse for interface all at the [edit protocols pim] hierarchy level on a PE router.
2. Configure PIM version 2 by including the version statement at the [edit protocols pim interface *interface-name*] hierarchy level.
3. Configure sparse mode (the mode in which the PIM interfaces operate) by including the mode statement at the [edit protocols pim interface *interface-name*] hierarchy level.
4. Configure the RP address by including the static statement at the [edit protocols pim rp] hierarchy level. In Figure 34 on page 417, the provider router is the RP.

Figure 34 on page 417 shows a multicast configuration on the provider network.

Figure 34: Multicast Configuration on the Provider Network

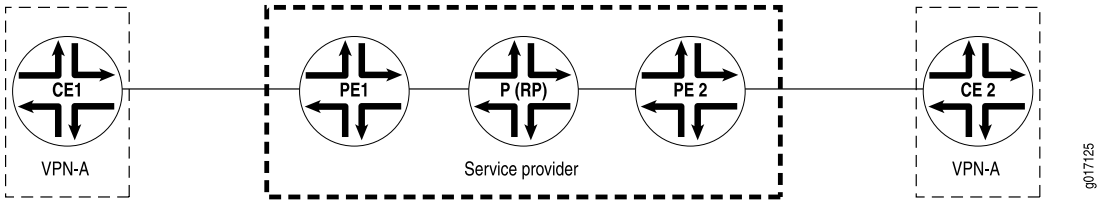

Configuring Multicast Connectivity on the CE Routers

To configure PIM for the master routing instance on all CE routers, include statements at the [edit protocols pim] hierarchy level:

1. Configure the interfaces going toward the provider router acting as the RP by including the interface statement at the [edit protocols pim] hierarchy level. In Figure 35 on page 418, the interfaces are labeled A and B.
2. Configure PIM version 2 by including the version statement at the [edit protocols pim interface *interface-name*] hierarchy level.
3. Configure sparse mode or sparse-dense mode (the mode in which the PIM interfaces operate) by including the mode statement at the [edit protocols pim interface *interface-name*] hierarchy level.
4. Configure the RP address by including the static statement at the [edit protocols pim rp] hierarchy level. In Figure 35 on page 418, CE2 is the RP router; however, the RP router can be anywhere in the customer network.

Figure 35 on page 418 shows multicast connectivity on the customer edge.

Figure 35: Multicast Connectivity on the CE Routers

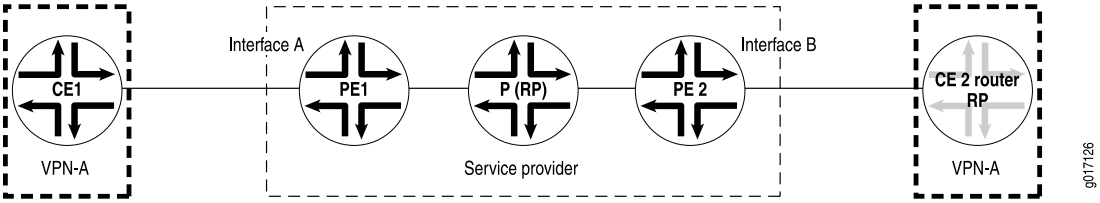

Configuring Multicast Connectivity for the VPN on the PE Router

To configure multicast connectivity for the VPN on the PE router, you must configure a VPN group address and configure the interfaces toward the router acting as RP. To configure the VPN group address, include the vpn-group-address statement at the [edit routing-instances *instance-name* protocols pim] hierarchy level:

```
[edit routing-instances instance-name protocols pim]
vpn-group-address address;
```

The PIM configuration in the VPN routing and forwarding (VRF) instance on the PE routers needs to match the master PIM instance on the CE router. Therefore, the PE router contains both a master PIM instance (to communicate with the provider core) and the VRF instance (to communicate with the CE routers). See the *JUNOS VPNs Configuration Guide* for information about configuring VPNs on PE routers.

NOTE: VRF instances that are part of the same VPN share the same VPN group address. For example, all PE routers containing multicast-enabled routing instance VPN-A share the same VPN group address configuration. In Figure 36 on page 418, the shared VPN group address configuration is 239.1.1.1.

Figure 36: Multicast Connectivity for the VPN

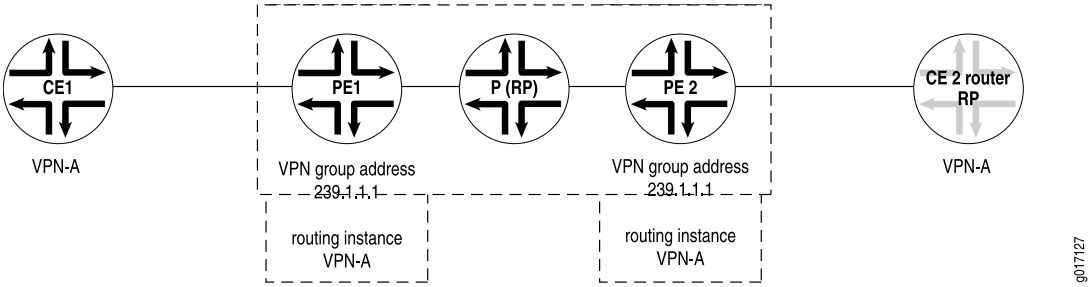

Configuring the Routing Group

Routing groups are usually configured at the [edit routing-instances *instance-name* routing-options] hierarchy level. However, with multicast in VRF instances, you must

configure routing groups differently. Configure the multicast routing group by adding the rib-groups statement at the [edit routing-options] hierarchy level.

After you configure the multicast routing group in the main routing instance, add the routing group to the VPN's VRF instance. To do this, include the rib-group statement at the [edit routing-instances *instance-name* protocols pim] hierarchy level.

For a multicast for Layer 3 VPN example, see “Example: Configuring PIM Sparse Mode over Layer 3 VPNs” on page 436.

Chapter 42

PIM Configuration Examples

This chapter includes the following Protocol Independent Multicast (PIM) configuration examples:

- Example: Configuring PIM Dense Mode on page 421
- Example: Configuring PIM Sparse-Dense Mode on page 422
- Example: Configuring PIM Sparse Mode on page 422
- Example: Configuring PIM Anycast with MSDP on page 423
- Example: Configuring PIM BSR Filters on page 427
- Example: Filtering Incoming PIM Join Messages on page 427
- Example: Configuring Register Message Filters on PIM RPs and DRs on page 428
- Example: Configuring Externally Facing PIM Border Routers on page 430
- Example: Tracing PIM Protocol Traffic on page 430
- Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 431
- Example: Configuring PIM Dense Mode over Layer 3 VPNs on page 431
- Example: Configuring PIM Sparse-Dense Mode over Layer 3 VPNs on page 433
- Example: Configuring PIM Sparse Mode over Layer 3 VPNs on page 436

Example: Configuring PIM Dense Mode

The following example shows a configuration for PIM dense mode:

```
protocols {
   pim {
      interface so-5/0/1 {
         mode dense;
      }
      interface so-5/0/2 {
         mode dense;
      }
      traceoptions {
         file log-pim;
         flag normal;
         flag state;
      }
   }
```

```
}
```

Example: Configuring PIM Sparse-Dense Mode

Configure PIM sparse-dense mode on all interfaces, specifying that the groups 224.0.1.39 and 224.0.1.40 are using dense mode:

```
protocols {
   pim {
      dense-groups {
         224.0.1.39;
         224.0.1.40;
      }
      interface all {
         version 1;
         mode sparse-dense;
      }
   }
}
```

Example: Configuring PIM Sparse Mode

This example shows how to configure a static RP router and non-RP routers for sparse mode.

1. Include the address statement at the [edit protocols pim rp local] hierarchy level.

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

NOTE: You do not need to configure Internet Group Management Protocol (IGMP) version 2 for a sparse mode configuration. After enabling PIM, by default, IGMP version 2 is also enabled.

```
protocols {
   pim {
      interface all {
         mode sparse;
         version 2;
      }
      interface fxp0.0 {
         disable;
      }
      rp {
         local {
            family inet {
               address 192.168.3.253;
            }
         }
```

```
        }
    }
}
```

2. In this example, you configure a non-RP router for PIM sparse mode. To specify a static RP address, include the address statement 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.

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

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

Example: Configuring PIM Anycast with MSDP

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 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 goes down, 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 covered in this section.

For information about standards supported for anycast RP, see “IP Multicast Specifications” on page 18.

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.

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 an network using an RP with MSDP, configure the RP using the lo0 or 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 198.58.3.254/32 and the shared RP address is 198.58.3.253/32. 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).

   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;
      }
   }
}
```

4. If you need to configure a PIM RP for both IPv4 and IPv6 scenarios, perform this step and the following one.

   Configure an RP using the lo0 or 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 198.58.3.254/32 and the shared RP address is 198.58.3.253/32. 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. Whether MSDP is used or not, the anycast RP configuration for a non-RP router is the same. 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 set PIM version 2.

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

7. Include the mode statement at the [edit protocols pim rp 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;
            }
        }
    }
```

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;
    }
}
```

Example: Filtering Incoming PIM Join Messages

In this example, you create the PIM join filter by including the **import pim-join-filter** statement at the [**edit protocols pim**] hierarchy level. Define **pim-join-filter** by adding the **policy-statement pim-join filter** statement at the [**edit policy-options**] hierarchy level. The filter is composed of a route filter and a source address filter—**bad-groups** and **bad-sources**, respectively. Policy **bad-groups** prevents (*,G) or (S,G) join messages from being received for all groups listed. Policy **bad-sources** 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.

```
protocols {
    pim {
        import pim-join-filter;
    }
}
policy-statement pim-join-filter {
    term bad-groups {
        from {
            route-filter 224.0.1.2/32 exact;
            route-filter 224.0.1.3/32 exact;
            route-filter 224.0.1.8/32 exact;
            route-filter 224.0.1.22/32 exact;
```

```
            route-filter 224.0.1.24/32 exact;
            route-filter 224.0.1.25/32 exact;
            route-filter 224.0.1.35/32 exact;
            route-filter 224.0.1.39/32 exact;
            route-filter 224.0.1.40/32 exact;
            route-filter 224.0.1.60/32 exact;
            route-filter 224.0.2.1/32 exact;
            route-filter 224.0.2.2/32 exact;
            route-filter 225.1.2.3/32 exact;
            route-filter 229.55.150.208/32 exact;
            route-filter 234.42.42.42/30 orlonger;
            route-filter 239.0.0.0/8 orlonger;
        }
        then reject;
    }
    term bad-sources {
        from {
            source-address-filter 10.0.0.0/8 orlonger;
            source-address-filter 127.0.0.0/8 orlonger;
            source-address-filter 172.16.0.0/12 orlonger;
            source-address-filter 192.168.0.0/16 orlonger;
        }
        then reject;
    }
    term last {
        then accept;
    }
}
```

Example: Configuring Register Message Filters on PIM RPs and DRs

Configure an RP filter to drop the register packets for multicast group range 224.1.1.0/24 from source address 10.10.94.1:

```
protocols {
    pim {
        rp {
            rp-register-policy incoming-policy-for-rp;
            local {
                address 10.10.10.5;
            }
        }
    }
}
policy-options {
    policy-statement incoming-policy-for-rp {
        from {
            route-filter 224.1.1.0/24 orlonger;
            source-address-filter 10.10.94.2/32 exact;
            then reject;
        }
    }
}
```

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:

```
protocols {
   pim {
      rp {
         dr-register-policy outgoing-policy-for-dr;
         static {
            address 10.10.10.3;
         }
      }
   }
}
policy-options {
   policy-statement outgoing-policy-for-rp {
      from {
         route-filter 224.1.1.0/24 orlonger;
         source-address-filter 10.10.10.1/32 exact;
         then reject;
      }
   }
}
```

More complex register message filtering is possible. This example configures a policy expression to accept register messages for multicast group 224.1.1.5 but reject those for 224.1.1.1:

```
protocols {
   pim {
      rp {
         rp-register-policy [ reject_224_1_1_1 | accept_224_1_1_5 ];
         local {
            address 10.10.10.5;
         }
      }
   }
}
policy-options {
   policy-statement reject_224_1_1_1 {
      from {
         route-filter 224.1.1.0/24 orlonger;
         source-address-filter 10.10.94.2/32 exact;
         then reject;
      }
      policy-statement accept_224_1_1_5 {
         term one {
            from {
               route-filter 224.1.1.5/32 exact;
               source-address-filter 10.10.94.2/32 exact;
            }
            then accept;
         }
         term two {
            then reject;
         }
      }
```

```
    }
}
```

Example: Configuring Externally Facing PIM Border Routers

In this example, you add the scope statement at the [edit routing-options multicast] hierarchy level to prevent auto-RP traffic from "leaking" into or out of your PIM domain. Two scopes defined below, auto-rp-39 and auto-rp-40, are for specific addresses. The scoped-range statement defines a group range, thus preventing group traffic from leaking.

```
routing-options {
    multicast {
        scope auto-rp-39 {
            prefix 224.0.1.39/32;
            interface t1-0/0/0.0;
        }
        scope auto-rp-40 {
            prefix 224.0.1.40/32;
            interface t1-0/0/0.0;
        }
        scope scoped-range {
            prefix 239.0.0.0/8;
            interface t1-0/0/0.0;
        }
    }
}
```

Example: Tracing PIM Protocol Traffic

Trace only unusual or abnormal operations to a routing log file, and trace detailed information about all PIM messages to a PIM log file:

```
routing-options {
    traceoptions {
        file routing-log;
        flag errors;
    }
}
protocols {
    pim {
        interface so-0/0/0;
        traceoptions {
            file pim-log;
            flag packets;
        }
    }
}
```

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 Dense Mode over Layer 3 VPNs

Multicast over Layer 3 VPNs for dense mode works much the same way as in sparse mode. In the following example, the VPN network uses dense mode for the entire multicast group range. Compare this with the configuration used in "Example: Configuring PIM Sparse Mode over Layer 3 VPNs" on page 436. In that configuration, sparse mode is used for the entire multicast group range.

1. Configure PIM on the P router.

```
protocols {
   pim {
      dense-groups {
         224.0.1.39/32;
         224.0.1.40/32;
      }
      rp {
         local {
```

```
            address 10.255.71.47;
         }
      }
      interface all {
         mode sparse;
         version 2;
      }
      interface fxp0.0 {
         disable;
      }
   }
}
```

2. Configure PIM on the PE router. Include the **mode** statement at the [**edit protocols pim interface**] hierarchy level to specify **sparse** mode.

```
protocols {
   pim {
      rp {
         static {
            address 10.255.71.47;
         }
      }
      interface all {
         mode sparse;
         version 2;
      }
      interface fxp0.0 {
         disable;
      }
   }
}
```

3. Configure PIM on the CE router. Include the **mode** statement at the [**edit protocols pim interface**] hierarchy level to specify **dense** mode.

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

4. Configure dense mode in the routing instance of the PE router facing the CE router. Include the **mode** statement at the [**edit routing-instances** *instance-name* **protocols pim**] hierarchy level to specify **dense** mode for the CE-facing interface **t1-1/0/0:0.0**.

```
routing-instances {
   VPN-A {
      instance-type vrf;
```

```
            interface t1-1/0/0:0.0;
            interface lo0.1;
            route-distinguisher 10.255.71.46:100;
            vrf-import VPNA-import;
            vrf-export VPNA-export;
            protocols {
                ospf {
                    export bgp-to-ospf;
                    area 0.0.0.0 {
                        interface t1-1/0/0:0.0;
                        interface lo0.1;
                    }
                }
                pim {
                    vpn-group-address 239.1.1.1;
                    interface t1-1/0/0:0.0 {
                        mode dense;
                        version 2;
                    }
                    interface lo0.1 {
                        mode dense;
                        version 2;
                    }
                }
            }
        }
    }
    interfaces {
        lo0 {
            description "unit 1 has the important PIM address";
            unit 0 {
                family inet {
                    address 192.168.27.13/32;
                    primary;
                    address 127.0.0.1/32;
                }
            }
            unit 1 {
                family inet {
                    address 10.10.47.101/32;
                }
            }
        }
    }
```

For information about configuring VPNs, see the *JUNOS VPNs Configuration Guide*.

Example: Configuring PIM Sparse-Dense Mode over Layer 3 VPNs

Multicast over Layer 3 VPNs for sparse-dense mode works much the same way as in sparse mode. In the following example, the VPN network uses dense mode for group range 229.0.0.0/8 and sparse mode for the remaining multicast group range outside 229.0.0.0/8. Compare this with the configuration used in “Example:

Configuring PIM Sparse Mode over Layer 3 VPNs" on page 436. In that configuration, sparse mode is used for the entire multicast group range.

1. Configure PIM on the P router.

```
protocols {
   pim {
      dense-groups {
         224.0.1.39/32;
         224.0.1.40/32;
      }
      rp {
         local {
            address 10.255.71.47;
         }
      }
      interface all {
         mode sparse;
         version 2;
      }
      interface fxp0.0 {
         disable;
      }
   }
}
```

2. Configure sparse-dense mode for the PE router interfaces by including the mode statement at the [edit protocols pim interface] hierarchy level.

```
protocols {
   pim {
      rp {
         static {
            address 10.255.71.47;
         }
      }
      interface all {
         mode sparse-dense;
         version 2;
      }
      interface fxp0.0 {
         disable;
      }
   }
}
```

3. Configure PIM on the CE router. Include the dense-groups statement at the [edit protocols pim] hierarchy level to define the desired dense groups address range on the CE router. Include the mode statement at the [edit protocols pim interface] hierarchy level to specify sparse-dense mode.

```
protocols {
   pim {
      dense-groups {
         229.0.0.0/8;
      }
```

```
        rp {
            static {
                address 10.255.245.91;
            }
        }
        interface all {
            mode sparse-dense;
            version 2;
        }
        interface fxp0.0 {
            disable;
        }
    }
}
```

4. Include the **dense-groups** statement at the [**edit routing-instances** *instance-name* **protocols pim**] hierarchy level to define the desired dense groups address range for the routing instance on the PE router. Include the **mode** statement at the [**edit routing-instances** *instance-name* **pim interface**] hierarchy level to specify **sparse-dense** mode for the CE-facing interface **t1-1/0/0:0.0**.

```
routing-instances {
    VPN-A {
        instance-type vrf;
        interface t1-1/0/0:0.0;
        interface lo0.1;
        route-distinguisher 10.255.71.46:100;
        vrf-import VPNA-import;
        vrf-export VPNA-export;
        protocols {
            ospf {
                export bgp-to-ospf;
                area 0.0.0.0 {
                    interface t1-1/0/0:0.0;
                    interface lo0.1;
                }
            }
            pim {
                dense-groups {
                    229.0.0.0/8;
                }
                vpn-group-address 239.1.1.1;
                rp {
                    static {
                        address 10.255.245.91;
                    }
                }
                interface t1-1/0/0:0.0 {
                    mode sparse-dense;
                    version 2;
                }
                interface lo0.1 {
                    mode sparse-dense;
                    version 2;
                }
            }
```

```
        }
      }
    }
    interfaces {
      lo0 {
        description "unit 1 has the important PIM address";
        unit 0 {
          family inet {
            address 192.168.27.13/32;
            primary;
            address 127.0.0.1/32;
          }
        }
        unit 1 {
          family inet {
            address 10.10.47.101/32;
          }
        }
      }
    }
```

For information about configuring VPNs, see the *JUNOS VPNs Configuration Guide*.

Example: Configuring PIM Sparse Mode over Layer 3 VPNs

This section describes how to configure multicast in PIM sparse mode for a range of multicast addresses for VPN-A as shown in Figure 37 on page 436.

Figure 37: Customer Edge and Service Provider Networks

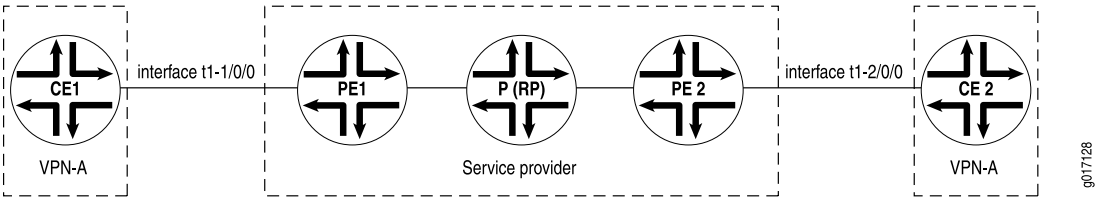

1. Configure PIM on the P router. The P router acts as the P (RP) router in this example. Specify the P router's address (10.255.71.47) at the [edit protocols pim rp local] hierarchy level.

```
protocols {
  pim {
    dense-groups {
      224.0.1.39/32;
      224.0.1.40/32;
    }
    rp {
      local {
        address 10.255.71.47;
      }
    }
```

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

2. Configure PIM on the provider edge 1 (PE1) router. Specify a static route to the service provider RP router—the P router (10.255.71.47).

```
protocols {
    pim {
        rp {
            static {
                address 10.255.71.47;
            }
        }
        interface all {
            mode sparse;
            version 2;
        }
        interface fxp0.0 {
            disable;
        }
    }
}
```

3. Configure PIM on the provider edge 2 (PE2) router. Specify a static route to the service provider RP—the P router (10.255.71.47).

```
protocols {
    pim {
        rp {
            static {
                address 10.255.71.47;
            }
        }
        interface all {
            mode sparse;
            version 2;
        }
        interface fxp0.0 {
            disable;
        }
    }
}
```

4. Configure PIM on the customer edge (CE1) router. Specify the RP address for the VPN RP—router CE2 (10.255.245.91).

```
protocols {
    pim {
        rp {
```

```
            static {
                address 10.255.245.91;
            }
        }
        interface all {
            mode sparse;
            version 2;
        }
        interface fxp0.0 {
            disable;
        }
    }
}
```

5. Configure PIM on the customer edge 2 (CE2) router, which acts as the VPN RP. Specify router CE2's address (10.255.245.91) at the [edit protocols pim rp local] hierarchy level:

```
protocols {
    pim {
        rp {
            local {
                address 10.255.245.91;
            }
        }
        interface all {
            mode sparse;
            version 2;
        }
        interface fxp0.0 {
            disable;
        }
    }
}
```

6. Configure the routing instance (VPN-A) for the Layer 3 VPN on router PE1. As part of the configuration, you need to establish the PIM instance for the VPN. Include the vpn-group-address statement at the [edit routing-instances *routing-instance-name* protocols pim] hierarchy level to specify the VPN group address, which is needed for multicast over a Layer 3 VPN configuration.

   Set the RP configuration for the VRF instance at the [edit routing-instances *routing-instance-name* protocols pim] hierarchy level. The RP configuration within the VRF instance provides explicit knowledge of the RP address, so that the (*,G) state can be forwarded.

   For Release 5.5 or later, configure an additional unit on the loopback interface of the PE router at the [edit interfaces] hierarchy level, and assign an address from the VPN address space. Then add the newly created loopback interface in two places:

   - Routing instance (VPN-A) at the [edit routing-instances *routing-instance-name*] hierarchy level.
   - Routing instance (VPN-A) at the [edit routing-instances *routing-instance-name* protocols pim] hierarchy level.

Also, add the loopback interface to the IGP and Border Gateway Protocol (BGP) policies to advertise the interface in the VPN address space. For more information about how to configure a logical unit on a loopback interface, see the *JUNOS VPNs Configuration Guide*.

In multicast Layer 3 VPNs, the multicast PE routers must use the primary loopback address (or router ID) for sessions with their internal BGP peers. If the PE routers use a route reflector with next-hop self configured, Layer 3 multicast over VPN will not work because PIM cannot transmit upstream interface information for multicast sources behind remote PEs into the network core. Multicast Layer 3 VPNs require that the BGP next-hop address of the VPN route match the BGP next-hop address of the loopback VRF instance address.

```
routing-instances {
   VPN-A {
      instance-type vrf;
      interface t1-1/0/0:0.0;
      interface lo0.1;
      route-distinguisher 10.255.71.46:100;
      vrf-import VPNA-import;
      vrf-export VPNA-export;
      protocols {
         ospf {
            export bgp-to-ospf;
            area 0.0.0.0 {
               interface t1-1/0/0:0.0;
               interface lo0.1;
            }
         }
         pim {
            vpn-group-address 239.1.1.1;
            rp {
               static {
                  address 10.255.245.91;
               }
            }
            interface t1-1/0/0:0.0 {
               mode sparse;
               version 2;
               interface lo0.1 {
                  mode sparse;
                  version 2;
               }
            }
         }
      }
   }
   interfaces {
      lo0 {
         description "unit 1 has the important PIM address";
         unit 0 {
            family inet {
               address 192.168.27.13/32;
               primary;
               address 127.0.0.1/32;
```

```
            }
        }
        unit 1 {
            family inet {
                address 10.10.47.101/32;
            }
        }
    }
  }
}
```

NOTE: Multicast Layer 3 VPNs require the BGP next-hop address of the VPN route to match the BGP next-hop address of the loopback VRF instance address.

7. Configure the routing instance (VPN-A) for the Layer 3 VPN on the PE2 router. You need to set the PIM instance for the VPN. Include the vpn-group-address statement at the [edit routing-instances *routing-instance-name* protocols pim] hierarchy level to specify the VPN group address, which is used for multicast over a Layer 3 VPN configuration. As you did for the PE1 router, configure an additional unit on the loopback interface of the PE2 router at the [edit interfaces] hierarchy level and assign an address from the VPN address space.

```
routing-instances {
    VPN-A {
        instance-type vrf;
        interface t1-2/0/0:0.0;
        interface lo0.1;
        route-distinguisher 10.255.71.51:100;
        vrf-import VPNA-import;
        vrf-export VPNA-export;
        protocols {
            ospf {
                export bgp-to-ospf;
                area 0.0.0.0 {
                    interface t1-2/0/0:0.0;
                    interface lo0.1;
                }
            }
            pim {
                vpn-group-address 239.1.1.1;
                rp {
                    static {
                        address 10.255.245.91;
                    }
                }
                interface t1-2/0/0:0.0 {
                    mode sparse;
                    version 2;
                    interface lo0.1 {
                        mode sparse;
                        version 2;
                    }
                }
```

```
            }
        }
    }
    interfaces {
        lo0 {
            description "unit 1 has the important PIM address";
            unit 0 {
                family inet {
                    address 192.168.27.14/32;
                    primary;
                    address 127.0.0.1/32;
                }
            }
            unit 1 {
                family inet {
                    address 10.10.47.102/32;
                }
            }
        }
    }
}
```

NOTE: Multicast Layer 3 VPNs require that the BGP next-hop address of the VPN route match the BGP next-hop address of the loopback VRF instance address.

8. When one of the PE routers is running Cisco Systems IOS software, you must configure the Juniper Networks PE router to support this multicast interoperability requirement. The Juniper Networks PE router must have the lo0.0 interface in the master routing instance and the lo0.1 interface assigned to the VPN routing instance. You must configure the lo0.1 interface with the same IP address that the lo0.0 interface uses for BGP peering in the provider core in the master routing instance.

   Configure the same IP address on the lo0.0 and lo0.1 loopback interfaces of the Juniper Networks PE router at the [edit interfaces lo0] hierarchy level, and assign the address used for BGP peering in the provider core in the master routing instance.

```
lo0 {
    description "unit 0 and unit 1 configured for Cisco IOS interoperability";
    unit 0 {
        family inet {
            address 192.168.27.14/32;
            primary;
            address 127.0.0.1/32;
        }
    }
    unit 1 {
        family inet {
            address 192.168.27.14/32;
        }
    }
}
```

9. Configure the multicast routing table group by including the VPNA-mcast-rib statement at the [edit routing-options] hierarchy level. This group accesses inet.2 when doing RPF checks. However, if you are using inet.0 for multicast RPF checks, this step will prevent your multicast configuration from working.

   You must also include the interface routes in inet.2. For more information about creating routing table groups, see the *JUNOS Routing Protocols Configuration Guide*.

```
routing-options {
   interface-routes {
      rib-group VPNA-mcast-rib;
   }
   rib-groups {
      VPNA-mcast-rib {
         export-rib VPN-A.inet.2;
         import-rib VPN-A.inet.2;
      }
   }
}
```

10. After you configure the multicast routing table group, activate it by including the rib-group inet VPNA-mcast-rib statement at the [edit routing-instances *instance-name* protocols pim] hierarchy level of the VPN's VRF instance.

```
routing-instances {
   VPN-A {
      protocols {
         pim {
            rib-group inet VPNA-mcast-rib;
         }
      }
   }
}
```

11. Display multicast tunnel information and the number of neighbors, using the show pim interfaces instance *instance-name* command from the PE1 or PE2 router. When issued from the PE1 router, the output display is:

```
user@host> show pim interfaces instance VPN-A

Instance: PIM.VPN-A
Name                  Stat Mode       IP V State Count DR address
lo0.1                 Up   Sparse      4 2 DR        0 10.10.47.101
mt-1/1/0.32769        Up   Sparse      4 2 DR        1
mt-1/1/0.49154        Up   Sparse      4 2 DR        0
pe-1/1/0.32769        Up   Sparse      4 1 P2P       0
t1-2/1/0:0.0          Up   Sparse      4 2 P2P       1
```

   You can also display all PE tunnel interfaces, using the show pim join command from the provider router acting as the RP.

12. Display multicast tunnel interface information, DR information, and the PIM neighbor status between VRF instances on PE1 and PE2, using the show pim

neighbors instance *instance-name* command from either PE router. When issued from the PE1 router, the output display is:

```
user@host> show pim neighbors instance VPN-A

Instance: PIM.VPN-A
Interface           IP V Mode        Option      Uptime Neighbor addr
mt-1/1/0.32769       4 2             HPL       01:40:46 10.10.47.102
t1-1/0/0:0.0         4 2             HPL       01:41:41 192.168.196.178
```

For information about configuring VPNs, see the *JUNOS VPNs Configuration Guide*.

Chapter 43

Summary of PIM Configuration Statements

The following sections explain each of the Protocol Independent Multicast (PIM) configuration statements. The statements are organized alphabetically.

address (Anycast RPs)

Syntax

address *address* <forward-msdp-sa>;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp local (inet | inet6) anycast-pim rp-set],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local (inet | inet6) anycast-pim rp-set],
[edit protocols pim rp local (inet | inet6) anycast-pim rp-set],
[edit routing-instances *routing-instance-name* protocols pim rp local (inet | inet6) anycast-pim rp-set]

Release Information

Statement introduced in JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure the anycast rendezvous point (RP) addresses in the RP set. Multiple addresses can be configured in an RP set. If the RP has peer Multicast Source Discovery Protocol (MSDP) connections, then the RP must forward MSDP source active (SA) messages.

Options

address—RP address in an RP set.

forward-msdp-sa—(Optional) Forward MSDP SAs to this address.

Required Privilege Level

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

address (Local RPs)

Syntax address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

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

Description Configure the local rendezvous point (RP) address.

Options *address*—Local RP address.

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

Related Topics
- Configuring the Local RP Address on page 398

address (Static RPs)

Syntax

```
address address {
   group-ranges {
      destination-ip-prefix</prefix-length>;
   }
   version version;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp static],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp static],
[edit protocols pim static],
[edit routing-instances *routing-instance-name* protocols pim rp static]

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure static rendezvous point (RP) addresses. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.

For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.

Options

address—Static RP address.
Default: 224.0.0.0/4

The remaining statements are explained separately.

Required Privilege Level

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

Related Topics

- Configuring Static PIM RPs on page 399

algorithm

Syntax algorithm *algorithm-name*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection authentication]
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection authentication]

Release Information Statement introduced in JUNOS Release 9.6.

Description Specify the algorithm to use for BFD authentication.

Options *algorithm-name*—Name of algorithm to use for BFD authentication:

- simple-password—Plain-text password. One to 16 bytes of plain text. One or more passwords may 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 Topics

- authentication statement
- Overview of BFD Authentication for PIM on page 379
- Configuring BFD Authentication for PIM on page 381

anycast-pim

Syntax

```
anycast-pim {
   rp-set {
      address address <forward-msdp-sa>;
   }
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information

Statement introduced in JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

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 Topics

- Example: Configuring PIM Anycast with MSDP on page 423

assert-timeout

Syntax assert-timeout *seconds*;

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

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

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 210 seconds. You can configure this assert timeout between 5 and 210 seconds.

Options *seconds*—Time for routing device to wait before another assert message cycle.
Range: 5 through 210 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 Topics
- Configuring the PIM Assert Timeout on page 413

authentication

Syntax

```
authentication {
   algorithm algorithm-name;
   key-chain key-chain-name;
   loose-check;
}
```

Hierarchy Level

[edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information

Statement introduced in JUNOS Release 9.6.

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 Topics

- bfd-liveness-detection statement
- key-chain statement
- loose-check statement
- Overview of BFD Authentication for PIM on page 379
- Configuring BFD Authentication for PIM on page 381
- Configuring the BFD Protocol for PIM on page 378

auto-rp

Syntax

```
auto-rp {
   (announce | discovery | mapping);
   (mapping-agent-election | no-mapping-agent-election);
}
```

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 Release 7.5.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure automatic RP announcement and discovery.

Options

announce—Configures the routing device to listen only for mapping packets and also to advertise itself if it is an RP.

discovery—Configures the routing device to listen only for mapping packets.

mapping—Configures the routing device to announce, listens for and generates mapping packets, and announces that the routing device is eligible to be an RP.

The remaining statement is explained separately.

Required Privilege Level

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

Related Topics

- Configuring PIM Auto-RP on page 402

bfd-liveness-detection

Syntax

```
bfd-liveness-detection {
   authentication {
      algorithm algorithm-name;
      key-chain key-chain-name;
      loose-check;
   }
   detection-time {
      threshold milliseconds;
   }
   minimum-interval milliseconds;
   minimum-receive-interval milliseconds;
   multiplier number;
   no-adaptation;
   transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
   }
   version (0 | 1 | automatic);
}
```

Hierarchy Level

[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information

Statement introduced in JUNOS Release 8.1.
authentication statement introduced in JUNOS Release 9.6.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure bidirectional forwarding detection (BFD) timers and authentication.

Required Privilege Level

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

Related Topics

- Configuring BFD Authentication for PIM on page 381
- Configuring the BFD Protocol for PIM on page 378

bootstrap

Syntax

```
bootstrap {
   family (inet | inet6) {
      export [ policy-names ];
      import [ policy-names ];
      priority number;
   }
}
```

Hierarchy Level

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

Release Information

Statement introduced in JUNOS Release 7.6.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

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 Topics

- Configuring the Routing Device's Bootstrap Router Priority on page 401

bootstrap-export

Syntax

bootstrap-export [*policy-names*];

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Apply one or more export policies to control outgoing PIM bootstrap messages.

Options

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

Required Privilege Level

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

Related Topics

- bootstrap-import
- Filtering PIM IPv4 Bootstrap Messages on page 400

bootstrap-import

Syntax bootstrap-import [*policy-names*];

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

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

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

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

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

Related Topics

- bootstrap-export
- Filtering PIM IPv4 Bootstrap Messages on page 400

bootstrap-priority

Syntax bootstrap-priority *number*;

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

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

Description Configure whether this routing device is eligible to be a bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router.

Options *number*—Priority for becoming the bootstrap router. A value of 0 means that the routing device is not eligible to be the bootstrap router.
Range: 0 through 255
Default: 0

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

Related Topics
- *JUNOS Multicast Protocols Configuration Guide*

dense-groups

Syntax
```
dense-groups {
   addresses;
}
```

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 Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure which groups are operating in dense mode.

Options *addresses*—Operate in dense mode.

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

Related Topics
- Configuring PIM Sparse-Dense Mode Properties on page 392

disable (PIM Graceful Restart)

Syntax disable;

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

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

Description Explicitly disable PIM sparse mode graceful restart.

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

Related Topics
- Configuring PIM Sparse Mode Graceful Restart on page 395

disable (PIM)

Syntax disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* protocols pim family (inet | inet6)],
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit protocols pim],
[edit protocols pim family (inet | inet6)],
[edit protocols pim interface *interface-name*],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim],
[edit routing-instances *routing-instance-name* protocols pim family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information
Statement introduced before JUNOS Release 7.4.
In JUNOS Release 9.6, the disable statement was extended to the [family] hierarchy level.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Explicitly disable PIM at the protocol, interface or family hierarchy levels.

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

Related Topics
- Disabling PIM on page 385
- family

dr-election-on-p2p

Syntax dr-election-on-p2p;

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

Release Information Statement introduced in JUNOS Release 9.1.
Statement introduced in JUNOS Release 9.1 for EX Series switches.

Description Enable PIM designated router (DR) election on point-to-point (P2P) links.

Default No PIM DR election is performed on P2P links.

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

Related Topics
- Configuring PIM Designated Router Election on Point-to-Point Links on page 373

dr-register-policy

Syntax dr-register-policy [*policy-names*];

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

Release Information Statement introduced in JUNOS Release 7.6.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

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

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

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

Related Topics
- rp-register-policy
- Configuring Register Message Filtering on a PIM RP or DR on page 407

embedded-rp

Syntax

```
embedded-rp {
   group-ranges {
      destination-ip-prefix</prefix-length>;
   }
   maximum-rps limit;
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

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 Topics

- Configuring PIM Embedded RP for IPv6 on page 412

export (Bootstrap)

Syntax

export [*policy-names*];

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp bootstrap family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap family (inet | inet6)],
[edit protocols pim rp bootstrap family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap family (inet | inet6)]

Release Information

Statement introduced in JUNOS Release 7.6.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Apply one or more export policies to control outgoing PIM bootstrap messages.

Options

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

Required Privilege Level

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

Related Topics

- import
- Filtering PIM Bootstrap Messages on page 402

export (PIM)

Syntax

export [*policy-names*];

Hierarchy Level

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

Release Information

Statement introduced in JUNOS Release 9.6.

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 Topics

- Filtering Outgoing PIM Join Messages on page 376

family (Bootstrap)

Syntax

```
family (inet | inet6) {
   export [ policy-names ];
   number;
   [ policy-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp bootstrap],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap],
[edit protocols pim rp bootstrap],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap]

Release Information

Statement introduced in JUNOS Release 7.6.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure which IP protocol type bootstrap properties to apply.

Options

inet—Apply IP version 4 (IPv4) local RP properties.

inet6—Apply IPv6 local RP properties.

The remaining statements are explained separately.

Required Privilege Level

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

Related Topics

- Configuring the Routing Device's Bootstrap Router Priority on page 401

family (Disable PIM)

Syntax

```
family (inet | inet6) {
    disable;
}
```

Hierarchy Level

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

Release Information

Statement introduced in JUNOS Release 9.6.

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 Topics

- disable
- disable
- Disabling PIM on page 385

family (Local RP)

Syntax

```
family (inet | inet6) {
   disable;
   address address;
   anycast-pim {
      local-address address;
      rp-set {
         address address <forward-msdp-sa>;
      }
   }
   group-ranges {
      destination-ip-prefix</prefix-length>;
   }
   hold-time seconds;
   priority number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp local],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local],
[edit protocols pim rp local],
[edit routing-instances *routing-instance-name* protocols pim rp local]

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure which IP protocol type local RP properties to apply.

Options

inet—Apply IP version 4 (IPv4) local RP properties.

inet6—Apply IPv6 local RP properties.

The remaining statements are explained separately.

Required Privilege Level

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

Related Topics

- *JUNOS Multicast Protocols Configuration Guide*

graceful-restart

Syntax

```
graceful-restart {
   disable;
   restart-duration seconds;
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure PIM sparse mode graceful restart.

The remaining statements are explained separately.

Required Privilege Level

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

Related Topics

- Configuring PIM Sparse Mode Graceful Restart on page 395

group-ranges

Syntax

```
group-ranges {
   destination-ip-prefix</prefix-length>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp embedded-rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp embedded-rp],
[edit protocols pim rp embedded-rp],
[edit protocols pim rp local family (inet | inet6)],
[edit protocols pim rp static address *address*],
[edit routing-instances *routing-instance-name* protocols pim rp embedded-rp],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp static address *address*]

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure the address ranges of the multicast groups for which this routing device can be an RP.

Default

The routing device is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12).

Options

destination-mask—Addresses or address ranges for which this routing device can be an RP.

Required Privilege Level

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

Related Topics

- Configuring the Groups for Which the Routing Device Is the RP on page 398
- Configuring PIM Embedded RP for IPv6 on page 412

hello-interval

Syntax hello-interval *seconds*;

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

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

Description Specify how often the router sends PIM hello packets out of an interface.

Options *seconds*—Length of time between PIM hello packets.
Range: 0 through 255
Default: 30 seconds

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

Related Topics
- hold-time
- Modifying the PIM Hello Interval on page 373

hold-time

Syntax hold-time *seconds*;

Hierarchy Level [edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

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

Description Specify the time period for which a neighbor is to consider the sending routing device (this routing device) to be operative (up).

Options *seconds*—Hold time.
Range: 0 through 255
Default: 0 seconds

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

Related Topics
- Modifying the Local RP Hold Time on page 399

import (Bootstrap)

Syntax import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp bootstrap (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)],
[edit protocols pim rp bootstrap (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)]

Release Information Statement introduced in JUNOS Release 7.6.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

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

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

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

Related Topics
- export
- Filtering PIM Bootstrap Messages on page 402

import (PIM)

Syntax import [*policy-names*];

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

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

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 from entering the network.

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

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

Related Topics
- Filtering Incoming PIM Join Messages on page 377

infinity

Syntax infinity [*policy-names*];

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

Release Information Statement introduced in JUNOS Release 8.0.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the infinity statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair.

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

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

Related Topics
- Configuring the PIM SPT Threshold Policy on page 413

interface

Syntax

```
interface (all | interface-name) {
   disable;
   bfd-liveness-detection {
      authentication {
         algorithmalgorithm-name;
         key-chainkey-chain-name;
         loose-check;
      }
      detection-time {
         threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      version (0 | 1 | automatic);
   }
   family (inet | inet6) {
      disable;
   }
   hello-interval seconds;
   mode (dense | sparse | sparse-dense);
   neighbor-policy [ policy-names ];
   priority number;
   version version;
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

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

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify all. For details about specifying interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

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 Topics

- *JUNOS Multicast Protocols Configuration Guide*

join-prune-timeout

Syntax join-prune-timeout *seconds*;

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

Release Information Statement introduced in JUNOS Release 8.4.

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 Topics
- Modifying the IGMP Join State Timeout on page 412

join-load-balance

Syntax join-load-balance;

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

Release Information Statement introduced in JUNOS Release 9.0.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Enable load balancing of PIM join messages across interfaces and routing devices.

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

Related Topics
- Configuring PIM Join Load Balancing on page 408

key-chain

Syntax key-chain *key-chain-name*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection authentication]
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection authentication]

Release Information Statement introduced in JUNOS Release 9.6.

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 Topics

- authentication
- Configuring BFD Authentication for PIM on page 381
- Overview of BFD Authentication for PIM on page 379

local

Syntax

```
local {
   disable;
   address address;
   family (inet | inet6) {
      disable;
      address address;
      anycast-pim {
         local-address address;
            rp-set {
            address address <forward-msdp-sa>;
         }
      }
      group-ranges {
         destination-ip-prefix</prefix-length>;
      }
      hold-time seconds;
      priority number;
   }
   group-ranges {
      destination-ip-prefix</prefix-length>;
   }
   hold-time seconds;
   priority number;
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.
The remaining statements are explained separately.

Description

Configure the routing device's RP properties.

Required Privilege Level

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

Related Topics

- Configuring Local PIM RPs on page 396

local-address

Syntax local-address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6) anycast-pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6) anycast-pim],
[edit protocols pim rp local family (inet | inet6) anycast-pim],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6) anycast-pim]

Release Information Statement introduced in JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the routing device's local address for anycast rendezvous point (RP). If this statement is omitted, the router ID is used as this address.

Options *address*—Anycast RP IPv4 or IPv6 address, depending on family configuration.

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

Related Topics
- Example: Configuring PIM Anycast with MSDP on page 423

loose-check

Syntax loose-check;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection authentication],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection authentication]

Release Information Statement introduced in JUNOS Release 9.6.

Description Optionally modify the default of strict checking for BFD authentication, and specify loose checking instead. When loose checking is configured, packets are accepted without authentication being checked at each end of the BFD session. The use of this statement is intended for transitional periods only.

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

Related Topics
- authentication
- Configuring BFD Authentication for PIM on page 381
- Overview of BFD Authentication for PIM on page 379

mapping-agent-election

Syntax (mapping-agent-election | no-mapping-agent-election);

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

Release Information Statement introduced in JUNOS Release 7.5.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the routing device's 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 Topics
- Configuring Auto-RP Mapping Agent Election on page 406

maximum-rps

Syntax maximum-rps *limit*;

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

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

Description Limit the number of RPs that the routing device acknowledges.

Options *limit*—Number of RPs.
Range: 1 through 500
Default: 100

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

Related Topics
- Configuring PIM Embedded RP for IPv6 on page 412

minimum-interval

Syntax minimum-interval *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in JUNOS Release 8.1.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the bidirectional forwarding detection (BFD) minimum interval timer. This timer specifies the same value for both the minimum transmit interval and minimum receive interval for the bfd-liveness-detection statement.

Options *milliseconds*—Minimum transmit and receive interval.
Range: 1 through 255,000 milliseconds

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

Related Topics
- Configuring the BFD Protocol for PIM on page 378

minimum-receive-interval

Syntax minimum-receive-interval *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in JUNOS Release 8.1.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the bidirectional forwarding detection (BFD) minimum receive interval timer. This timer specifies only the minimum receive interval for the bfd-liveness-detection statement.

Options *milliseconds*—Minimum receive interval.
Range: 1 through 255,000 milliseconds

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

Related Topics
- Configuring the BFD Protocol for PIM on page 378

mode

Syntax mode (dense | sparse | sparse-dense);

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

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

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.

Related Topics

- Configuring PIM Dense Mode Properties on page 391
- Configuring PIM Sparse-Dense Mode Properties on page 392
- *JUNOS Multicast Protocols Configuration Guide*

multiplier

Syntax multiplier *number*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in JUNOS Release 8.1.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the multiplier for bidirectional forwarding detection (BFD) timers.

Options *number*—Detection time multiplier.
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 Topics
- Configuring the BFD Protocol for PIM on page 378

neighbor-policy

Syntax neighbor-policy [*policy-names*];

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

Release Information Statement introduced in JUNOS Release 8.2.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Apply a PIM interface-level policy to filter neighbor IP addresses.

Options *policy-name*—Name of the policy that filters neighbor IP addresses. For details about configuring policy statements, see the *JUNOS Policy Framework Configuration Guide*.

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

Related Topics
- Configuring Interface-Level PIM Neighbor Policies on page 373.

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 {
      disable;
      bfd-liveness-detection {
         authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
         }
         detection-time {
            threshold milliseconds;
         }
         minimum-interval milliseconds;
         minimum-receive-interval milliseconds;
         multiplier number;
         version (0 | 1 | automatic);
      }
      family (inet | inet6) {
         disable;
      }
      hello-interval seconds;
      mode (dense | sparse | sparse-dense);
      neighbor-policy [ policy-names ];
      priority number;
      version version;
   }
   join-load-balance;
   join-prune-timeout;
   nonstop-routing;
   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 {
                    rp-set {
                        address address <forward-msdp-sa>;
                    }
                    disable;
                    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 {
                version version;
                group-ranges {
                    destination-ip-prefix</prefix-length>;
                }
            }
        }
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

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

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

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.

Related Topics
- Configuring PIM Dense Mode Properties on page 391
- Configuring PIM Dense Mode Properties on page 391
- *JUNOS Multicast Protocols Configuration Guide*

priority (Bootstrap)

Syntax priority *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp bootstrap (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)],
[edit protocols pim rp bootstrap (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)]

Release Information Statement introduced in JUNOS Release 7.6.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the routing device's likelihood to be elected as the bootstrap router.

Options *number*—Routing device's priority for becoming the bootstrap router. A higher value corresponds to a higher priority.
Range: 0 through a 32-bit number
Default: 0 (The routing device has the least likelihood of becoming the bootstrap router and sends packets with a priority of 0.)

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

Related Topics
- bootstrap-priority
- Configuring the Routing Device's Bootstrap Router Priority on page 401

priority (PIM Interfaces)

Syntax priority *number*;

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

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

Description Configure the routing device's likelihood to be elected as the designated router.

Options *number*—Routing device's priority for becoming the designated router. A higher value corresponds to a higher priority.
Range: 1 through a 32-bit number
Default: 1 (The routing device has the least likelihood of becoming the designated router.)

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

Related Topics
- Configuring Interface Priority to Become the PIM Designated Router on page 372

priority (PIM RPs)

Syntax priority *number*;

Hierarchy Level [edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

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

Description Configure this routing device's priority for becoming an RP. The bootstrap router uses this field when selecting the list of candidate RPs to send in the bootstrap message. A smaller number increases the likelihood that the routing device becomes the RP for local multicast groups. A priority value of 0 means that bootstrap router can override the group range being advertised by the candidate RP.

Options *number*—Routing device's priority for becoming an RP. A lower value corresponds to a higher priority.
Range: 0 through 255
Default: 1

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

Related Topics
- Configuring Local PIM RPs on page 396

restart-duration

Syntax restart-duration *seconds*;

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

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

Description Configure the duration of the graceful restart interval.

Options *seconds*—Time the routing device waits (in seconds) to complete PIM sparse mode graceful restart.
Range: 30 through 300
Default: 60

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

Related Topics
- Configuring PIM Sparse Mode Graceful Restart on page 395

rib-group

Syntax rib-group *group-name*;

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

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

Description Associate a routing table group with PIM.

Options *group-name*—Name of the routing table group. The name must be one that you defined with the rib-group 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 Topics
- Configuring a PIM RPF Routing Table on page 374

rp

Syntax

```
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-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;
   }
   local {
      family (inet | inet6) {
         disable;
         address address;
         anycast-pim {
            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 ];
   static {
      address address {
         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 Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

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 Topics
- *JUNOS Multicast Protocols Configuration Guide*

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 Release 7.6.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

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 Topics
- dr-register-policy
- Configuring Register Message Filtering on a PIM RP or DR on page 407

rp-set

Syntax

```
rp-set {
   address address <forward-msdp-sa>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim local family (inet | inet6) anycast-pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim local family (inet | inet6) anycast-pim],
[edit protocols pim local family (inet | inet6) anycast-pim],
[edit routing-instances *routing-instance-name* protocols pim local family (inet | inet6) anycast-pim]

Release Information

Statement introduced in JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs.

The remaining statements are explained separately.

Required Privilege Level

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

Related Topics

- Example: Configuring PIM Anycast with MSDP on page 423

spt-threshold

Syntax

```
spt-threshold {
   infinity [ policy-names ];
}
```

Hierarchy Level

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

Release Information

Statement introduced in JUNOS Release 8.0.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Set the SPT threshold to infinity for a source-group address pair. Last-hop multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or an SPT rooted at the source. By default, last-hop routing devices transition to a direct SPT to the source. You can configure this routing device to set the SPT transition value to infinity to prevent this transition for any source-group address pair.

The remaining statements are explained separately.

Required Privilege Level

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

Related Topics

- Configuring the PIM SPT Threshold Policy on page 413

static

Syntax

```
static {
   address address {
      group-ranges {
         destination-ip-prefix</prefix-length>;
      }
      version version;
   }
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure static RP addresses. The default static RP address is 224.0.0.0/4. To configure other addresses, include one or more address statements. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.

For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.

The remaining statements are explained separately.

Required Privilege Level

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

Related Topics

- Configuring Static PIM RPs on page 399

traceoptions

Syntax

```
traceoptions {
   file filename <files number> <size size> <world-readable | no-world-readable>;
   flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description

Configure PIM tracing options.

To specify more than one tracing operation, include multiple flag statements.

Default

The default PIM trace options are those inherited from the routing protocol's traceoptions statement included at the [edit routing-options] hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place tracing output in the pim-log file.

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.
Range: 2 through 1000 files
Default: 2 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.

PIM Tracing Flags

- assert—Assert messages
- bootstrap—Bootstrap messages
- cache—Packets in the PIM sparse mode routing cache

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

  Default: If you do not specify this option, only unusual or abnormal operations are traced.
- policy—Policy operations and actions
- route—Routing table changes
- state—State transitions
- task—Interface transactions and processing
- timer—Timer usage

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

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

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.
Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

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

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

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

If you specify a maximum file size, you must also include the files statement to specify the maximum number of trace files.
Syntax: *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB
Range: 0 KB through the maximum file size supported on your system
Default: 1 MB

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

Required Privilege Level

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

Related Topics

- Configuring PIM Trace Options on page 388
- Tracing DVMRP Protocol Traffic on page 324
- Tracing MSDP Protocol Traffic on page 512
- Configuring PIM Trace Options on page 388

version (BFD)

Syntax version (0 | 1 | automatic);

Hierarchy Level [edit protocols piminterface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in JUNOS Release 8.1.
Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Specify the bidirectional forwarding detection (BFD) protocol version that you want to detect.

Options 0—BFD version 0; this version is deprecated.

1—BFD version 1.

automatic—Auto-detect the BFD version (default).

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

Related Topics

- Configuring the BFD Protocol for PIM on page 378

version (PIM)

Syntax version *version*;

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

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

Description Specify the version of PIM.

Options *version*—PIM version number.
Range: 1 or 2
Default: PIM version 2

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

Related Topics
- Changing the PIM Version on page 372

vpn-group-address

Syntax vpn-group-address *address*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify a group address on which to encapsulate multicast traffic from a virtual private network (VPN) instance.

Options *address*—IP address 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. For more information about addresses, see "Multicast Addresses" on page 20.

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

Related Topics
- Configuring Any-Source Multicast for Draft-Rosen VPNs on page 415

Part 12
MSDP

- MSDP Introduction on page 495
- MSDP Configuration Guidelines on page 497
- Summary of MSDP Configuration Statements on page 515

Chapter 44

MSDP Introduction

MSDP Overview

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 Border Gateway Protocol (BGP). 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 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 this peer is toward 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 "Configuring MSDP Mesh Groups" on page 502.

Router R determines 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 Border Gateway Protocol (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.

For more information about PIM sparse mode, see *JUNOS Multicast Protocols Configuration Guide*.

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.

For information about standards supported for MSDP, see the “IP Multicast Specifications” on page 18.

Chapter 45

MSDP Configuration Guidelines

This chapter describes the following tasks for configuring MSDP:

- Configuring MSDP on page 497
- Minimum MSDP Configuration on page 499
- Enabling MSDP on page 499
- Configuring MSDP Peers on page 500
- Configuring MSDP Groups on page 501
- Configuring MSDP Mesh Groups on page 502
- Configuring the MSDP Authentication Key on page 503
- Configuring MSDP Routing Policy on page 504
- Configuring Multiple Rendezvous Points in a Domain on page 506
- Configuring MSDP Data Encapsulation on page 507
- Configuring the MSDP Active Source Limit on page 508
- Configuring a Default MSDP Peer on page 510
- Disabling MSDP on page 511
- Tracing MSDP Protocol Traffic on page 512
- Example: Configuring MSDP on page 513

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 ...
    }
    import [ policy-names ];
```

```
    local-address address;
    peer address {
        ... peer-configuration ...
    }
    rib-group group-name;
    source ip-prefix</prefix-length> {
        active-source-limit {
            maximum number;
            threshold number;
        }
    }
    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]

For an overview of logical systems and a detailed example of logical system configuration, see the logical systems chapter of the *JUNOS Feature Guide*.

By default, MSDP is disabled.

For a configuration example, see "Example: Configuring MSDP" on page 513.

Minimum MSDP Configuration

To enable MSDP on the router, include at least the following statements:

```
msdp {
    local-address address;
    peer address;
}
```

You can include these statements 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]

You must configure at least one peer. The peer and the local-address statements are required. The router must also be a Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP). For configuration information, see "Configuring Local PIM RPs" on page 396.

Enabling MSDP

To enable MSDP peering on the router, include the msdp statement:

```
msdp {
    local-address address;
    peer address;
    rib-group group-name;
}
```

You can include this statement at the following hierarchy levels:

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

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

To associate with MSDP a routing table group that imports and exports routes into the specified routing table group, include the rib-group statement. The routing table group is a group that you defined with the rib-groups statement at the [edit routing-options] hierarchy level. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring MSDP Peers

An MSDP router must know which routers are its peers. You define the peer relationships explicitly by configuring the neighboring routers that are the MSDP peers of the local router. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. You must configure at least one peer for MSDP to function.

To configure MSDP peers, include the peer statement:

```
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>;
    }
}
```

The peer and the local-address statements are required.

You can configure MSDP peers globally or in a group.

- Globally for all MSDP peers at the following hierarchy levels:
  - [edit protocols msdp]
  - [edit logical-systems *logical-system-name* protocols msdp]
  - [edit routing-instances *routing-instance-name* protocols msdp]
  - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp]
- In a group at the following hierarchy levels:

- [edit protocols group *group-name*]
- [edit logical-systems *logical-system-name* protocols group *group-name*]
- [edit routing-instances *routing-instance-name* protocols group *group-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols group *group-name*]

If you configure MSDP peers in a group, each individual peer in a group inherits all group-level options.

Configuring MSDP Groups

You can arrange MSDP peers into groups. Each group must contain at least one peer. Arranging peers into groups is useful if you want to block sources from some peers and accept them from others, or set tracing options on one group and not others.

To configure MSDP groups, include one or more of the following statements:

```
group group-name {
   disable;
   export [ policy-names ];
   import [ policy-names ];
   local-address address;
   mode (mesh-group | standard);
   traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
   }
   peer address; {
      disable;
      active-source-limit {
         maximum number;
         threshold number;
      }
      authentication-key peer-key;
      default-peer;
      export [ policy-names ];
      import [ policy-names ];
      local-address address;
      traceoptions {
         file filename <files number> <size size> <world-readable | no-world-readable>;
         flag flag <flag-modifier> <disable>;
      }
   }
}
```

You can include these statements at the following hierarchy levels:

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

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp]
- [edit routing-instances *routing-instance-name* protocols msdp]

The local-address and peer statements are mandatory. The individual statements are discussed in separate sections.

Configuring MSDP Mesh Groups

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.

To configure an MSDP mesh group, define a peer group, and include the mode mesh-group statement:

```
group group-name {
    local-address address;
    mode mesh-group;
    peer address;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols msdp]
- [edit logical-systems *logical-system-name* protocols msdp]
- [edit routing-instances *routing-instance-name* protocols msdp]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp]

CAUTION: When configuring MSDP mesh groups, you must configure all members the same. 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 via 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.

Table 10 on page 503 explains how flooding is handled by peers in this configuration. Figure 38 on page 503 illustrates source-active message flooding between different mesh groups and peers within the same mesh group.

Table 10: 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 38: Source-Active Message Flooding

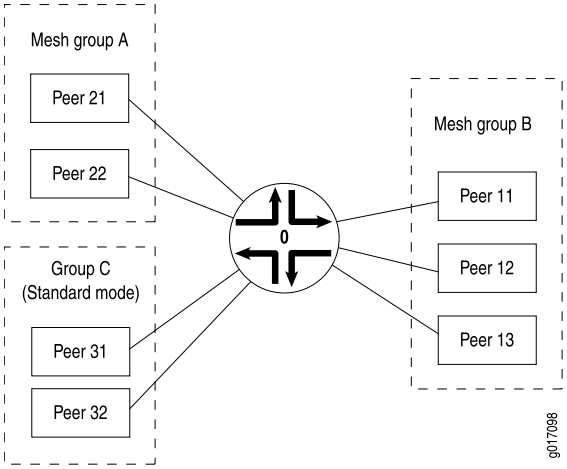

Configuring the MSDP Authentication Key

By default, multicast routers accept and process any properly formatted MSDP messages from the configured peer address. This default behavior might violate the security policies in many organizations because MSDP messages by definition come from another routing domain beyond the control of the security practices of the multicast router's organization.

The router can authenticate MSDP messages using the TCP message digest 5 (MD5) signature option for MSDP peering sessions. This authentication provides protection against spoofed packets being introduced into an MSDP peering session. Two organizations implementing MSDP authentication must decide on a human-readable key on both peers. This key is included in the MD5 signature computation for each MSDP segment sent between the two peers.

You configure an MSDP authentication key on a per-peer basis, whether the MSDP peer is defined in a group or individually. If you configure different authentication keys for the same peer at the [edit protocols msdp] and [edit protocols msdp group] hierarchy levels, the authentication key configured at the [edit protocols msdp] hierarchy level is used.

To configure MSDP authentication keys on the router, include the authentication-key statement:

```
authentication-key peer-key;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

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 (“ ”).

The following example configures the MSDP authentication key grandmother for MSDP peer 10.0.0.1, and the MSDP authentication keys New York and phoenix5 for peers 172.16.0.1 and 192.168.0.1 in MSDP group msdp-one:

```
msdp {
   group msdp-one {
      peer 171.16.0.1 {
         authentication-key "New York";
         local-address 10.100.0.2;
      }
      peer 192.168.0.1 {
         authentication-key phoenix5;
         local-address 10.100.0.2;
      }
      peer 10.0.0.1 {
         authentication-key grandmother;
         local-address 10.100.0.2;
      }
   }
}
```

Adding, removing, or changing an MSDP authentication key in a peering session resets the existing MSDP session and establishes a new session between the affected MSDP peers. This immediate session termination prevents excessive retransmissions and eventual session timeouts due to mismatched keys.

Configuring MSDP Routing Policy

All routing protocols use the routing table to store the routes that they learn and to determine which routes they advertise in their protocol packets. Routing policy allows you to control which routes the routing protocols store in, and retrieve from, the routing table. For information about routing policy, see the *JUNOS Routing Protocols Configuration Guide*.

You can configure routing policy globally, for a group, or for an individual peer:

- Globally for all MSDP peers at the [edit protocols msdp] or [edit logical-systems *logical-system-name* protocols msdp] hierarchy level.
- For all peers in a group at the [edit protocols msdp group *group-name*] or [edit logical-systems *logical-system-name* protocols msdp group *group-name*] level.
- For an individual peer at the [edit protocols msdp peer *address*] or [edit logical-systems *logical-system-name* protocols msdp peer *address*] level, or the [edit protocols msdp group *group-name* peer *address*] or [edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*] level.

If you configure routing policy at the group level, each individual peer in a group inherits the group's routing policy.

To apply policies to source-active messages being imported into the source-active cache from MSDP, include the import statement, listing the names of one or more policy filters to be evaluated. See Table 11 on page 505 for a list of match conditions.

Table 11: MSDP Source-Active Message 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 source-active message)
route-filter	Multicast group address embedded in the source-active message
source-address-filter	Multicast source address embedded in the source-active message

If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found, MSDP shares with the routing table only those routes that were learned from MSDP routers.

import [*policy-names*];

To apply policies to source-active messages being exported from the source-active cache into MSDP, include the export statement, listing the names of one or more policies to be evaluated. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the source-active cache entry. If no match is found, the default MSDP export policy is applied to entries in the source-active cache.

export [*policy-names*];

Configuring Multiple Rendezvous Points in a Domain

You can configure multiple RPs in a shared-tree PIM sparse-mode domain. You need to configure an MSDP local address to enable the RPs in the domain to maintain a consistent view of the active sources.

To configure a router to act as an RP in a domain with other RPs, do the following for each router in the domain that acts as an RP:

- Create the router ID by configuring a unique and routable IP address on the loopback interface and setting the primary address flag.
- Configure a non-unique, but routable, unicast address on the loopback interface.
- Use the non-unique, routable unicast address to configure the PIM router to be the local RP.
- Configure MSDP with the unique and routable address (router ID) as the local address of the peer.

For a sample configuration of multiple RPs, see "Example: Configuring a Router to Use Anycast RP" on page 506. For more information about configuring interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

Example: Configuring a Router to Use Anycast RP

The following example configures a router to use anycast RP:

```
[edit]
interfaces {
   ...
   lo0 {
      unit 0 {
         family inet {
            # This must be a unique routable address and router-id
            address 10.1.1.1/32 {
               primary;
            }
            # This must be a non-unique routable anycast RP address
            address 10.10.10.10/32;
            address 127.0.0.1/32;
         }
      }
   }
}
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;
        }
    }
    autonomous-system 1234;
}
protocols {
    bgp {
        group red {
            type internal;
            family inet any;
            neighbor 10.1.1.2 {
                local-address 10.1.1.1;
            }
        }
    }
    msdp {
        rib-group mcrg;
        group red {
            peer 10.1.1.2 {
                local-address 10.1.1.1;
            }
        }
    }
    pim {
        dense-groups {
            224.0.1.39/32;
            224.0.1.40/32;
        }
        rib-group mcrg;
        rp {
            local {
                address 10.10.10.10;
            }
        }
        interface all {
            mode sparse-dense;
            version 2;
        }
        interface fxp0.0 {
            disable;
        }
    }
}
```

Configuring MSDP Data Encapsulation

MSDP data encapsulation mainly concerns bursty sources of multicast traffic. Sources that send only one packet every few minutes have problems 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. So 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 denial-of-service (DoS) attack, the router's forwarding table may 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.

To configure MSDP data encapsulation on the router, include the data-encapsulation statement:

```
data-encapsulation (enable | disable);
```

You can include this statement at the following hierarchy levels:

- [edit protocols msdp]
- [edit logical-systems *logical-system-name* protocols msdp]
- [edit routing-instances *routing-instance-name* protocols msdp]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp]

The router must also be a PIM sparse-mode RP. For configuration information, see “Configuring Local PIM RPs” on page 396.

Configuring the MSDP Active Source Limit

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 DoS attack on a router running MSDP. To minimize this possibility, you can configure the router to limit the number of source active messages the router accepts. Also, you can configure a threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.

By default, the router accepts 25,000 source active messages before ignoring the rest to prevent a possible DoS attack. 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 1,000 messages are screened by the RED profile and the accepted messages processed.

To configure the MSDP active source limit on the router, include the active-source-limit statement:

```
active-source-limit {
   maximum number;
   threshold number;
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

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 at several levels of the MSDP hierarchy:

- Configuring Global, Group, and Peer Active Source Limit on page 509
- Configuring Per-Source Active Source Limit on page 510

Configuring Global, Group, and Peer Active Source Limit

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, all are applied.

The following example applies a limit of 5,000 active sources to MSDP peer 10.0.0.1, a limit of 7,500 active sources to MSDP peer 10.10.10.10 in group MSDP-group, and a limit of 10,000 active sources to all others.

```
[edit protocols msdp]
active-source-limit {
   maximum 10000;
}
group MSDP-group {
   peer 10.10.10.10;
   active-source-limit {
      maximum 7500;
   }
   peer 10.10.10.11;
}
peer 10.0.0.1 {
   active-source-limit {
      maximum 5000;
   }
}
```

Configuring Per-Source Active Source Limit

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.

```
[edit protocols msdp]
source 10.1.1.1/32 {
   active-source-limit {
      maximum 10000;
   }
}
source 10.1.0.0/16 {
   active-source-limit {
      maximum 500;
   }
}
source 0.0.0.0/0 {
   active-source-limit {
      maximum 5;
   }
}
```

In this example, the source 10.1.1.1 is allowed active sources for 10,000 groups. Any other source on the 10.1.0.0/16 network is allowed 500 groups. All other sources are allowed to source 5 active streams.

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 5,000 (and there are no other sources or limits configured), only 5,000 active source messages are accepted from this source.

Configuring a Default MSDP Peer

When a source-active message is received, a peer-RPF check is performed to make sure the peer is leading toward the originating RP and to decide whether to accept the source-active message. However, 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.

To establish an MSDP peer as the default peer, include the **default-peer** statement:

```
default-peer;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can establish a default peer at the peer or group level.

Related Topics

- MSDP Overview on page 495

Disabling MSDP

To disable MSDP on the router, include the **disable** statement:

```
disable;
```

You can disable MSDP globally for all peers, for all peers in a group, or for an individual peer.

- Globally for all MSDP peers at the following hierarchy levels:
  - [edit protocols msdp]
  - [edit logical-systems *logical-system-name* protocols msdp]
  - [edit routing-instances *routing-instance-name* protocols msdp]
  - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp]
- For all peers in a group at the following hierarchy levels:
  - [edit protocols msdp group *group-name*]
  - [edit logical-systems *logical-system-name* protocols msdp group *group-name*]
  - [edit routing-instances *routing-instance-name* protocols msdp group *group-name*]
  - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*]
- For an individual peer at the following hierarchy levels:
  - [edit protocols msdp peer *address*]
  - [edit protocols msdp group *group-name* peer *address*]
  - [edit logical-systems *logical-system-name* protocols msdp peer *address*]
  - [edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*]
  - [edit routing-instances *routing-instance-name* protocols msdp peer *address*]

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*]

If you disable MSDP at the group level, each peer in the group is disabled.

Tracing MSDP Protocol Traffic

To trace MSDP protocol traffic, you can specify options in the global traceoptions statement at the [edit routing-options] or [edit logical-systems *logical-system-name* routing-options] hierarchy level, and you can specify MSDP-specific options by including the traceoptions statement. Options applied at the routing options level trace all packets, and options applied at the protocol level trace only IGMP traffic.

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

You can configure MSDP tracing as follows:

- Globally for all MSDP peers at the [edit protocols msdp] hierarchy level
- For all peers in a group at the [edit protocols msdp group *group-name*] hierarchy level
- For an individual peer at the [edit protocols msdp peer *address*] or the [edit protocols msdp group *group-name* peer *address*] hierarchy level

If you configure tracing options at the group level, each peer in the group inherits the group's tracing options.

You can configure tracing options globally for all MSDP peers (at the [edit protocols msdp] hierarchy level), for all peers in a group (at the [edit protocols msdp group *group-name*] level), or for an individual peer (at the [edit protocols msdp peer *address*] or the [edit protocols msdp group *group-name* peer *address*] level). If you configure tracing options at the group level, each peer in the group inherits the group's tracing options.

You can specify the following MSDP-specific options in the flag statement:

- keepalive—Trace keepalive messages.
- packets—Trace all MSDP packets.
- route—Trace MSDP changes to the routing table.
- sa—Trace source-active packets.
- sa-request—Trace source-active request packets.
- sa-response—Trace source-active response packets.

For general information about tracing, see the *JUNOS System Basics Configuration Guide*.

Example: Tracing MSDP Protocol Traffic

Trace only unusual or abnormal operations to routing-log, and trace detailed information about all MSDP messages to msdp-log:

```
[edit]
routing-options {
   traceoptions {
      file routing-log;
      flag errors;
   }
}
protocols {
   msdp {
      peer 192.68.2.120; {
         local-address 192.68.1.200;
      }
      traceoptions {
         file msdp-log;
         flag packets;
      }
   }
}
```

Example: Configuring MSDP

Configure a router to act as a PIM sparse-mode rendezvous point and an MSDP peer:

```
[edit]
routing-options {
   interface-routes {
      rib-group ifrg;
   }
   rib-groups {
      ifrg {
         import-rib [inet.0 inet.2];
      }
      mcrg {
         export-rib inet.2;
         import-rib inet.2;
      }
   }
}
protocols {
   bgp {
      group lab {
         type internal;
         family any;
         neighbor 192.168.6.18 {
            local-address 192.168.6.17;
         }
```

```
        }
    }
    pim {
        dense-groups {
            224.0.1.39/32;
            224.0.1.40/32;
        }
        rib-group mcrg;
        rp {
            local {
                address 192.168.1.1;
            }
        }
        interface all {
            mode sparse-dense;
            version 1;
        }
    }
    msdp {
        rib-group mcrg;
        group lab {
            peer 192.168.6.18 {
                local-address 192.168.6.17;
            }
        }
    }
}
```

Chapter 46

Summary of MSDP Configuration Statements

The following sections explain each of the Multicast Source Discovery Protocol (MSDP) configuration statements. The statements are organized alphabetically.

active-source-limit

Syntax

```
active-source-limit {
    maximum number;
    threshold number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* protocols msdp source *ip-address/prefix-length*],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp source *ip-address/prefix-length*],
[edit protocols msdp],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit protocols msdp source *ip-address/prefix-length*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp source *ip-address/prefix-length*]

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Limit the number of active source messages the router 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.

Usage Guidelines

See “Configuring the MSDP Active Source Limit” on page 508.

Required Privilege Level

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

authentication-key

Syntax authentication-key *peer-key*;

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

Release Information Statement introduced before JUNOS Release 7.4.

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 router accepts any valid MSDP messages from the peer address.

Options *peer-key*—MD5 authentication key. The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of (,), &, and [. If you include spaces in an MSDP authentication key, enclose all characters in quotation marks (“ ”).

Usage Guidelines See “Configuring the MSDP Authentication Key” on page 503.

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

data-encapsulation

Syntax data-encapsulation (disable | enable);

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a rendezvous point (RP) using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages.

Default If you do not include this statement, the RP encapsulates multicast data.

Options disable—(Optional) Do not use MSDP data encapsulation.

enable—Use MSDP data encapsulation.
Default: enable

Usage Guidelines See "Configuring MSDP Data Encapsulation" on page 507.

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

default-peer

Syntax default-peer;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Establishes this peer as the default MSDP peer and accepts source-active messages from the peer without the usual peer-reverse-path-forwarding (peer-RPF) check.

Usage Guidelines See “Configuring a Default MSDP Peer” on page 510.

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

disable

Syntax disable;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Explicitly disable MSDP.

Usage Guidelines See “Disabling MSDP” on page 511.

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

export

Syntax export [*policy-names*];

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Apply one or more policies to routes being exported from the routing table into MSDP.

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

Usage Guidelines See “Configuring MSDP Routing Policy” on page 504.

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

Related Topics import

group

Syntax

```
group group-name {
   disable;
   export [ policy-names ];
   import [ policy-names ];
   local-address address;
   mode (mesh-group | standard);
   traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
   }
   peer address; {
      disable;
      active-source-limit {
         maximum number;
         threshold number;
      }
      authentication-key peer-key;
      default-peer;
      export [ policy-names ];
      import [ policy-names ];
      local-address address;
      traceoptions {
         file filename <files number> <size size> <world-readable | no-world-readable>;
         flag flag <flag-modifier> <disable>;
      }
   }
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Define an MSDP peer group. MSDP peers within groups share common tracing options, if present and not overridden for an individual peer with the peer statement. To configure multiple MSDP groups, include multiple group statements.

By default, the group's options are identical to the global MSDP options. To override the global options, include group-specific options within the group statement.

The group must contain at least one peer.

Options

group-name—Name of the MSDP group.

The remaining statements are explained separately.

Usage Guidelines

See “Configuring MSDP Routing Policy” on page 504.

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

import

Syntax import [*policy-names*];

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Apply one or more policies to routes being imported into the routing table from MSDP.

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

Usage Guidelines See “Configuring MSDP Routing Policy” on page 504.

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

Related Topics export

local-address

Syntax local-address *address*;

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

Release Information Statement introduced before JUNOS Release 7.4.

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.

Usage Guidelines See “Enabling MSDP” on page 499, “Configuring MSDP Peers” on page 500, and “Configuring MSDP Groups” on page 501.

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

maximum

Syntax maximum *number*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the maximum number of MSDP active source messages the router accepts.

Options *number*—Maximum number of active source messages.
Range: 1 through 1,000,000
Default: 25,000

Usage Guidelines See "Configuring the MSDP Active Source Limit" on page 508.

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

Related Topics
- threshold

mode

Syntax mode (mesh-group | standard);

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

Release Information Statement introduced before JUNOS Release 7.4.

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—(Optional) Group of peers that are mesh group members.

standard—Use standard MSDP source-active flooding rules.
Default: standard

Usage Guidelines See "Configuring MSDP Mesh Groups" on page 502.

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

msdp

Syntax

```
msdp {
   disable;
   active-source-limit {
      maximum number;
      threshold number;
   }
   data-encapsulation (disable | enable);
   export [ policy-names ];
   group group-name {
      ... group-configuration ...
   }
   import [ policy-names ];
   local-address address;
   peer address {
      ... peer-configuration ...
   }
   rib-group group-name;
   source ip-prefix</prefix-length> {
      active-source-limit {
         maximum number;
         threshold number;
      }
   }
   traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
   }
   group group-name {
      disable;
      export [ policy-names ];
      import [ policy-names ];
      local-address address;
      mode <mesh-group | standard>;
      peer address {
         ... same statements as at the  [edit protocols msdp peer address] hierarchy level
            shown just following ...
      }
      traceoptions {
         file filename <files number> <size size> <world-readable | no-world-readable>;
         flag flag <flag-modifier> <disable>;
      }
   }
   peer address {
      disable;
      active-source-limit {
         maximum number;
         threshold number;
      }
      authentication-key peer-key;
      default-peer;
      export [ policy-names ];
```

```
        import [ policy-names ];
        local-address address;
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
    }
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Enable MSDP on the router. You must also configure at least one peer for MSDP to function.

Default

MSDP is disabled on the router.

Options

The statements are explained separately.

Usage Guidelines

See "Enabling MSDP" on page 499.

Required Privilege Level

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

peer

Syntax

```
peer address {
   disable;
   active-source-limit {
      maximum number;
      threshold number;
   }
   authentication-key peer-key;
   default-peer;
   export [ policy-names ];
   import [ policy-names ];
   local-address address;
   traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
   }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*]

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Define an MSDP peering relationship. An MSDP router must know which routers are its peers. You define the peer relationships explicitly by configuring the neighboring routers that are the MSDP peers of the local router. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. To configure multiple MSDP peers, include multiple peer statements.

By default, the peer's options are identical to the global or group-level MSDP options. To override the global or group-level options, include peer-specific options within the peer statement.

At least one peer must be configured for MSDP to function. You must configure *address* and local-address.

Options

address—Name of the MSDP peer.

The remaining statements are explained separately.

Usage Guidelines

See “Configuring MSDP Peers” on page 500.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

rib-group

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 Release 7.4.

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.

Usage Guidelines See "Enabling MSDP" on page 499.

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

source

Syntax

```
source ip-address</prefix-length> {
   active-source-limit {
      maximum number;
      threshold number;
   }
}
```

Hierarchy Level

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

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Limit the number of active source messages the router accepts from sources in this address range.

Default

If you do not include this statement, the router accepts any number of MSDP active source messages.

Options

The other statements are explained separately.

Usage Guidelines

See “Configuring Per-Source Active Source Limit” on page 510.

Required Privilege Level

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

threshold

Syntax threshold *number*;

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

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the random early discard (RED) threshold for MSDP active source messages. This number must be less than the configured or default maximum.

Options *number*—RED threshold for active source messages.
Range: 1 through 1,000,000
Default: 24,000

Usage Guidelines See "Configuring the MSDP Active Source Limit" on page 508.

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

Related Topics
- maximum

traceoptions

Syntax

```
traceoptions {
   file filename <files number> <size size> <world-readable | no-world-readable>;
   flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Configure MSDP tracing options.

To specify more than one tracing operation, include multiple flag statements.

Default

The default MSDP trace options are those inherited from the routing protocol's traceoptions statement included at the [edit routing-options] hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place tracing output in the msdp-log file.

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.
Range: 2 through 1000 files
Default: 2 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.

MSDP Tracing Flags

- keepalive—Keepalive messages
- packets—All MSDP packets
- route—MSDP changes to the routing table
- source-active—Source-active packets
- source-active-request—Source-active request packets
- source-active-response—Source-active response packets

Global Tracing Flags

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

  Default: If you do not specify this option, only unusual or abnormal operations are traced.
- policy—Policy operations and actions
- route—Routing table changes
- state—State transitions
- task—Interface transactions and processing
- timer—Timer usage

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

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

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.
Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

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

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

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

If you specify a maximum file size, you must also include the files statement to specify the maximum number of trace files.
Syntax: *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB
Range: 10 KB through the maximum file size supported on your system
Default: 1 MB

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

Usage Guidelines

See "Tracing MSDP Protocol Traffic" on page 512.

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.

Part 13
Index

- Index on page 539
- Index of Statements and Commands on page 551

Index

Symbols

#, comments in configuration statements.....xxxvi
( ), in syntax descriptions.....xxxvi
< >, in syntax descriptions.....xxxv
[ ], in configuration statements.....xxxvi
{ }, in configuration statements.....xxxvi
| (pipe), in syntax descriptions.....xxxvi

A

accounting statement
    IGMP.....75
    IGMP interface.....75
        usage guidelines.....70
    IGMP protocol
        usage guidelines.....70
    MLD.....117
    MLD interface.....117
        usage guidelines.....113
    MLD protocol
        usage guidelines.....112
active-source-limit statement.....516
    usage guidelines.....508
address statement
    anycast RPs.....445
        usage guidelines.....423
    local RPs.....446
        usage guidelines.....398
    static RPs.....447
        usage guidelines.....399
addresses
    multicast.....8
administrative scoping.....9
algorithm statement
    BFD authentication.....448
all (tracing flag)
    PGM.....158
anycast RP.....506
anycast-pim statement.....449
    usage guidelines.....423
asm-override-ssm statement.....245
assert (tracing flag).....488
assert timeout
    configuring.....413
assert-timeout statement.....450
    usage guidelines.....413
authentication configuration
    BFD.....381
authentication statement
    BFD.....453
    BFD protocol.....451
authentication-key statement
    MSDP.....517
        usage guidelines.....503
auto-rp statement.....452
    usage guidelines.....402
autodiscovery.....166
    configuring.....173
autodiscovery statement.....177
autodiscovery-only statement.....178

B

backup-pe-group statement.....246
    usage guidelines.....235
backups statement.....247
    usage guidelines.....235
bandwidth statement.....248
    usage guidelines.....218
BFD
    authentication configuration.....381
    protocol.....378
BFD authentication
    algorithm statement.....448
    authentication statement.....451
    key-chain statement.....471
    loose-check statement.....473
bfd-liveness-detection statement
    PIM.....453
        usage guidelines.....378
bootstrap (tracing flag).....488
bootstrap IPv4 messages.....400
bootstrap messages.....346, 400, 401, 402
bootstrap statement.....454
bootstrap-export statement.....454
    usage guidelines.....400
bootstrap-import statement.....455
    usage guidelines.....400
bootstrap-priority statement.....456
    usage guidelines.....400, 401

braces, in configuration statements.....xxxvi
brackets
angle, in syntax descriptions.....xxxv
square, in configuration statements.....xxxvi
bridge domains
and IGMP snooping.....292
BSR
policy, import.....467

C

cache (tracing flag).....488
CBT
defined.....16
issues.....27
comments, in configuration statements.....xxxvi
configuration statements
hierarchy
multicast.....33
conventions
text and syntax.....xxxv
Core Based Trees *See* CBT
curly braces, in configuration statements.....xxxvi
customer support.....xxxvi
contacting JTAC.....xxxvi

D

data multicast distribution trees *See* MDT
data-encapsulation statement.....518
usage guidelines.....507
data-mdt-reuse statement.....179
default multicast distribution trees *See* MDT
default-peer statement.....519
usage guidelines.....510
dense-groups statement.....456
usage guidelines.....392
designated router.....344
disable statement
DVMRP.....329
usage guidelines.....323
IGMP.....76
usage guidelines.....73
MLD.....118
usage guidelines.....115
MSDP.....520
usage guidelines.....511
PIM family.....458
PIM graceful restart.....457
usage guidelines.....395
PIM interfaces.....458
PIM protocol.....458
SAP.....143
SAP and SDP
usage guidelines.....141
Distance Vector Multicast Routing Protocol *See* DVMRP
distribution trees
RPT.....12
shared trees.....12
documentation
comments on.....xxxvi
dr-election-on-p2p statement.....459
PIM
usage guidelines.....373
dr-register-policy statement.....459
usage guidelines.....407
draft-rosen MVPNs.....165, 166, 169
configuration statements.....177
DVMRP
configuration statements.....319
defined.....15
disabling.....323
enabling.....321, 330
hold-time period.....322
metric.....322, 333
overview.....317
policy, routing.....323, 331, 332
routing tables.....321, 334
dvmrp statement.....330
usage guidelines.....321

E

embedded RP
IPv6
configuring.....412
overview.....366
embedded-rp statement.....460
usage guidelines.....412
enable IGMP static group membership.....63
enable MLD static group membership.....105
event recording
IGMP.....69
MLD.....112
example
ingress PE redundancy.....237
exclude statement
IGMP.....76
MLD.....118
export statement
DVMRP.....331
usage guidelines.....323
MSDP.....521
usage guidelines.....504
PIM.....460, 461
configuring.....376
PIM RP
usage guidelines.....402

F

family statement
bootstrap....461
local RP....463
PIM
usage guidelines....397
PIM interfaces....462
PIM protocol....462
flood groups
and multicast snooping....279
flood-groups statement
multicast snooping....283
flow-map statement....249
usage guidelines....211
font conventions....xxxv
forwarding cache
and multicast snooping....278
forwarding cache limits, overview....229
forwarding class....227
forwarding-cache statement
multicast snooping....284
usage guidelines....229, 230
frames
multicast snooping....21, 23

G

graceful restart
and multicast snooping....279
disabling....395
overview....20
PIM
sparse mode....395
sparse-dense mode....392
graceful-restart statement
multicast snooping....285
PIM....464
usage guidelines....395
snooping
usage guidelines....279
graft (tracing flag)
DVMRP....335
PIM....488
group
provider tunnel....173
group range
MDT....170
group statement
IGMP....77
configuring....63
IGMP snooping....299
MDT....179
usage guidelines....171
MLD....119
configuring....105
MSDP....522
usage guidelines....501
group statement, IGMP snooping
IGMP snooping
usage guidelines....297
group-address statement....180
usage guidelines....173
group-count statement
IGMP....77
configuring....63
MLD....120
configuring....105
group-increment statement
IGMP....78
configuring....63
MLD....120
configuring....105
group-limit statement
IGMP snooping....300
group-limit statement, IGMP snooping
IGMP snooping
usage guidelines....297
group-policy statement
IGMP....78
usage guidelines....59
MLD....121
usage guidelines....103
group-range statement....180
usage guidelines....170
group-ranges statement....465
usage guidelines....398, 399
groups
SSM....266

H

hello (tracing flag)
PIM....488
hello-interval statement
PIM....466
usage guidelines....373
hold-time statement
DVMRP....331
usage guidelines....322
PIM....467
usage guidelines....399
host-only-interface statement....301
host-only-interface statement, IGMP snooping
IGMP snooping
usage guidelines....297

I

icons defined, notice....xxxiv

IGMP
and nonstop active routing....73
configuration statements....55, 56, 75, 293
configuring....55, 293
configuring PIM-to-IGMP message translation....241
disabling....73
enabling....57, 79
event recording....69
host-query message interval....57, 85
last-member query interval....60, 86
overview....53
PIM-to-IGMP message translation....239
query response interval....57, 86
robustness variable....61, 87
snooping and bridge domains....292
snooping interfaces....290
snooping overview....289
snooping proxies....291
static group membership....63
tracing operations....71
version....62, 92
IGMP snooping
enabling....302
group limit....300
host-only interface....301
host-query message interval....307
last-member query interval....308
multicast-router interface....305
proxy....306
query response interval....309
robust count....310
source address....311
igmp statement....79
usage guidelines....57
igmp-snooping statement....302
IGMPv3....53
interoperability with older versions....53
ignore-stp-topology-change statement....285
usage guidelines....280
immediate-leave statement
IGMP....80
usage guidelines....58
IGMP snooping....303
MLD....122
usage guidelines....102
immediate-leave statement, IGMP snooping
usage guidelines....294
import statement
bootstrap....467
usage guidelines....402
DVMRP....332
usage guidelines....323
MSDP....523
usage guidelines....504
PIM....468
usage guidelines....377
inclusive statement....181
inet-mdt statement
autodiscovery....181
infinity statement....468
usage guidelines....413
ingress PE redundancy
configuring....235
example....237
overview....235
init (tracing flag)
PGM....158
interface lists....9
interface statement
DVMRP....332
usage guidelines....321
IGMP....81
usage guidelines....57
IGMP snooping....304
MLD....123
usage guidelines....100
multicast....252
multicast scoping....253
usage guidelines....194
PIM....469
usage guidelines....391
interface statement,
IGMP snooping
usage guidelines....297
interfaces, multicast
pd....362
pe....362
pimd....362
pime....362
Internet Group Management Protocol *See* IGMP
interoperability
PIM....436
intra-as statement....182
IPv6
embedded RP
configuring....412
overview....366

J

join (tracing flag)....488
join-load-balance statement....470
usage guidelines....408
join-prune-timeout statement....470

K

keepalive (tracing flag)
MSDP....533
key-chain statement
BFD authentication....471

L

Layer 3 VPNs....163, 169, 415
leave (tracing flag)
IGMP....90
MLD....133
listen statement....144
usage guidelines....141
load balancing
for PIM join....408
local statement
PIM....472
usage guidelines....396
local-address statement....253
MSDP
usage guidelines....499
MSDP group....524
usage guidelines....501
MSDP peer....524
usage guidelines....500
PIM....473
usage guidelines....235
loose-check statement
BFD authentication....473

M

manuals
comments on....xxxvi
mapping-agent-election statement....474
usage guidelines....406
mappings
SSM....267
maximum statement....525
usage guidelines....508
maximum-bandwidth statement....254
usage guidelines....218
maximum-rps statement....474
usage guidelines....412
maximum-transmit-rate statement
IGMP....82
configuring....62
MLD....124
configuring....105
MBone....19
MDT....164
configuring....169
data-mdt-reuse....179
group address....179
group range....170, 180
rate limit....186
source address....188
statements....183
threshold....171
threshold parameters....189
tunnel characteristics....165
tunnel limit....172, 190
mdt statement....183
usage guidelines....169
mesh groups
MSDP....502
metric statement
DVMRP....333
usage guidelines....322
metrics
DVMRP....322, 333
minimum-interval statement
PIM....475
usage guidelines....378
minimum-receive-interval statement
PIM....453, 475
usage guidelines....378
MLD
configuration statements....100, 117
configuring PIM-to-MLD message translation....242
disabling....115
enabling....125
event recording....112
host-query message interval....101, 128
immediate-leave host removal
configuring....102
example....103
last-member query interval....102, 128
overview....95
PIM-to-MLD message translation....239
query response interval....101, 129
robustness variable....104, 129
static group membership....105
tracing operations....114
mld statement....125
usage guidelines....99
mode statement
DVMRP....333
usage guidelines....324
MSDP....526
usage guidelines....502
PIM....476
usage guidelines....392, 394
MOSPF, defined....15
MSDP
active source limit....516
maximum....525
per-source....531
threshold....532
authentication....503, 517

configuration statements.....497, 515
configuring.....497
data-encapsulation.....518
default peer.....510, 519
disabling.....511
enabling.....527
groups.....501, 522
local address.....524
mode.....526
peers.....500
policy, routing.....521, 523
routing tables.....530
tracing operations.....512
msdp statement.....527
usage guidelines.....499
mt (tracing flag).....488
mtrace (tracing flag)
IGMP.....71
MLD.....114
multicast
addresses.....8
backbone.....19
building blocks.....8
configuration statements.....255
configuring PIM-to-IGMP message translation.....241
configuring PIM-to-MLD message translation.....242
defined.....3
forwarding cache limits.....229
ingress PE redundancy.....235
Layer 2 frames.....21, 23
Layer 3 VPNs.....163, 169, 358, 415
leaf and branch.....7
MBone.....19
packet replication.....21
PIM-to-IGMP and PIM-to-MLD message translation.....239
protocols.....7, 14
group membership.....14
routing options.....193
configuration statements.....245
routing protocols.....15
compared, table.....16
scoping.....193, 194, 264
snooping.....21, 23
snooping and flood groups.....279
snooping and forwarding cache.....278
snooping and graceful restart.....279
snooping configuration statements.....277
snooping overview.....275
SSM groups.....266
SSM mapping.....267
standards supported.....17, 18
terminology.....6
uses.....5
Multicast Backbone *See* MBone
multicast distribution trees *See* MDT
multicast egress class of service.....227
multicast filters.....363
MAC filters.....364
MSDP SA messages.....365
RP/DR register messages.....364
configuring.....407
Multicast Listener Discovery *See* MLD
Multicast Open Shortest Path First *See* MOSPF
Multicast Source Discovery Protocol *See* MSDP
multicast statement.....255
usage guidelines.....194, 218, 229, 230
multicast-router-interface statement.....305
multicast-router-interface statement, IGMP snooping
IGMP snooping
usage guidelines.....297
multicast-snooping-options statement.....286
multiplier statement
PIM.....453, 477
usage guidelines.....378
mvpn statement.....184, 185

N

neighbor (tracing flag).....335
neighbor-policy statement
PIM.....477
usage guidelines.....373
neighbors
MSDP.....500
no-accounting statement
IGMP.....75
usage guidelines.....70
MLD.....117
usage guidelines.....113
no-qos-adjust statement.....257
nonstop active routing
role in PIM.....384
role of IGMP.....73
notice icons defined.....xxxiv
nsr-synchronization (tracing flag).....489

O

oif-map statement
IGMP (interface).....83
MLD (interface).....127

P

packets (tracing flag)
DVMRP.....335
IGMP.....90
MLD.....133

PGM....158
PIM....489
parentheses, in syntax descriptions....xxxvi
parser (tracing flag)
PGM....158
passive statement
IGMP (interface)....84
MLD (interface)....127
peer statement
MSDP....529
usage guidelines....500
PGM
architecture....150
configuration statements....157
configuring....155
overview....149
routers....152
tracing operations....158
pgm statement....157
usage guidelines....155
PIM
and nonstop active routing....384
anycast RP....449, 485
assert timeout....450, 486
configuring....413
background....27
BFD....378, 453, 475, 477, 491
bootstrap messages import and export....400, 402
bootstrap router....346, 400, 401
configuration statements....445, 492
configuring....369, 393
configuring PIM messages to IGMP messages....241
configuring PIM messages to MLD messages....242
dense mode....29, 355, 391
defined....15
designated router....344
embedded RP....460
configuring....412
overview....366
enabling....478
filters *See* multicast filters
graceful restart
disabling....395
overview....20
sparse mode....395
sparse-dense mode....392
hello interval....373
hold-time period....331, 467
incoming join filter policy, applying....377
interfaces
pd....362
pe....362
pimd....362
pime....362
interoperability....436
join load balancing
configuring....408
join-prune-timeout....470
Layer 3 VPNs....358
maximum RPs....474
mixing modes....31
modes....28
network components....28
outgoing join filter policy, applying....376
overview....341
policy, routing....468
rendezvous point tree....346
restart-duration statement....482
usage guidelines....395
routing tables....482
RPs....12, 344, 393, 396, 403, 483
anycast....449
anycast RP....357
embedded....460
mapping options....344
maximum....474
source registration....347
SPT cutover....350
SPT cutover control....354
sparse mode....29, 394
defined....16
sparse-dense mode....28, 357, 456
defined....15
SSM....30, 204, 354
translating PIM messages to IGMP and MLD messages....239
version....372, 492
VPN group address....418
pim statement....478
usage guidelines....369, 393
PIM-RP
SPT
configuring threshold cutover policy....413
pim-ssm statement....185
pim-to-igmp-proxy statement....258
pim-to-mld-proxy statement....259
policy statement
flow map....260
SSM map....261
policy, import
BSR....467
policy, routing
DVMRP....323, 331, 332
MSDP....504, 521, 523
PIM....468
PIM join filter....376, 377
Pragmatic General Multicast *See* PGM
prefix statement....261
usage guidelines....194

priority statement
bootstrap....480
PIM....481
usage guidelines....372, 398
probe (tracing flag)....335
promiscuous-mode statement....85
usage guidelines....60
Protocol Independent Multicast *See* PIM
protocols
group membership....14
multicast....7, 14
multicast routing....15
compared, table....16
provider tunnel
group address....173
provider-tunnel statement
usage guidelines....169
proxy statement
IGMP snooping....306
proxy statement, IGMP snooping
usage guidelines....294
prune (tracing flag)
DVMRP....335
PIM....489

Q

query-interval statement
IGMP....85
usage guidelines....57
IGMP snooping....307
MLD....128
usage guidelines....101
query-interval statement, IGMP snooping
usage guidelines....295
query-last-member-interval statement
IGMP....86
usage guidelines....60
IGMP snooping....308
MLD....128
usage guidelines....102
query-last-member-interval statement, IGMP snooping
usage guidelines....295
query-response-interval statement
IGMP....86
usage guidelines....57
IGMP snooping....309
MLD....129
usage guidelines....101
query-response-interval statement, IGMP snooping
usage guidelines....295

R

rate statement
MDT....186
usage guidelines....171
redundant-sources statement....262
register (tracing flag)....489
rendezvous point *See* PIM
rendezvous-point tree *See* RPT
replication
multicast packet....21
report (tracing flag)
DVMRP....335
IGMP....91
MLD....134
restart-duration statement....482
PIM graceful restart
usage guidelines....395
reverse path forwarding *See* RPF
reverse-oif-mapping statement....263
usage guidelines....226
rib-group statement
DVMRP....334
usage guidelines....321
MSDP....530
usage guidelines....499
PIM....482
usage guidelines....391
robust-count statement
IGMP....87
usage guidelines....61
IGMP snooping....310
MLD....129
usage guidelines....104
robust-count statement, IGMP snooping
usage guidelines....295
route (tracing flag)
MSDP....533
route-socket (tracing flag)
PGM....158
routing options
multicast....193
configuration statements....245
routing tables
DVMRP....321, 334
MSDP....530
PIM....482
routing-instances statement
vpn-group-address
usage guidelines....418
RP
anycast....449
embedded....460
rp (tracing flag)....489
rp statement....483
usage guidelines....393

rp-register-policy statement.....484
usage guidelines.....407
rp-set statement.....485
usage guidelines.....423
RPF.....10
checks.....13
policies.....199
table.....13
populating.....14
RPF check, multicast
RPF policy.....264
rpf-check-policy statement.....264
usage guidelines.....200
RPs
maximum.....474
PIM.....393
See also PIM RPs
RPT.....12

S

SAP
configuration statements.....143
configuring.....141
overview.....139
sap statement.....145
usage guidelines.....141
scope statement.....264
usage guidelines.....194
scope-policy statement.....265
usage guidelines.....195
scoping, multicast.....193, 264
with scope policy.....265
SDP.....139
Session Announcement Protocol *See* SAP
Session Description Protocol *See* SDP
shared trees.....12
shortest-path tree *See* SPT
show (tracing flag)
PGM.....158
signaling statement.....187
single forwarder election
configuring.....175
snooping *See* multicast overview
configuration statements.....277
flood groups and 279
forwarding cache and 278
graceful restart and 279
IGMP and VLANs.....298
IGMP interfaces.....290
IGMP overview.....289
IGMP proxies.....291
IGMP tracing operations.....298
multicast.....21, 23
spanning tree interfaces state changes.....280
source filtering.....53
source statement
IGMP.....87
configuring.....63
IGMP snooping.....310
MDT.....188
usage guidelines.....171
MLD.....130
configuring.....105
MSDP.....531
usage guidelines.....510
SSM.....266
usage guidelines.....208, 212
source-active (tracing flag).....533
source-active-request (tracing flag).....533
source-active-response (tracing flag).....533
source-address statement
IGMP snooping.....311
source-address statement, IGMP snooping
usage guidelines.....294
source-count statement
IGMP.....88
configuring.....63
MLD.....130
configuring.....105
source-increment statement
IGMP.....88
configuring.....63
MLD.....131
configuring.....105
source-specific multicast *See* SSM
SPT.....11
configuring threshold cutover policy.....413
cutover control.....350, 354
spt-threshold statement.....486
usage guidelines.....413
SSM.....30, 204
configuring.....206
domains.....206
mapping.....208
ssm-groups statement.....266
usage guidelines.....204
ssm-map statement
IGMP.....89
usage guidelines.....208
MLD.....131
usage guidelines.....208
SSM.....267
usage guidelines.....208
state (tracing flag)
PGM.....158
static statement
IGMP.....89
configuring.....63
IGMP snooping.....312

MLD....132
configuring....105
PIM....487
usage guidelines....399
static statement, IGMP snooping
IGMP snooping
usage guidelines....297
subscriber-leave-timer statement....268
usage guidelines....226
support, technical *See* technical support
syntax conventions....xxxv

T

technical support
contacting JTAC....xxxvi
threshold
MDT....171
threshold statement
forwarding cache....269
usage guidelines....229
MDT....189
usage guidelines....171
MSDP....532
usage guidelines....508
multicast snooping....287
timeout statement
flow map....270
forwarding cache....271
traceoptions statement
DVMRP....335
usage guidelines....324
IGMP....90
usage guidelines....71
MLD....133
usage guidelines....114
MSDP....533
usage guidelines....512
PGM....158
usage guidelines....155
PIM....488
usage guidelines....388
traceoptions statement, IGMP snooping
usage guidelines....298
tracing flags
all
PGM....158
assert....488
bootstrap....488
cache, PIM....488
graft
DVMRP....335
PIM....488
hello
PIM....488
init
PGM....158
join....488
keepalive
MSDP....533
leave
IGMP....90
MLD....133
MLD
leave....133
mt....488
mtrace
IGMP....71
MLD....114
neighbor....335
nsr-synchronization....489
packets
DVMRP....335
IGMP....90
MLD....133
PGM....158
PIM....489
parser, PGM....158
probe....335
prune
DVMRP....335
PIM....489
register....489
report
DVMRP....335
IGMP....91
MLD....134
route
MSDP....533
route-socket
PGM....158
rp....489
show
PGM....158
source-active....533
source-active-request....533
source-active-response....533
state
PGM....158
tracing operations
DVMRP....324, 335
IGMP....71, 90
IGMP snooping....298
MLD....114, 133
MSDP....512, 533
PGM....158
PIM....488
tunnel limit
MDT....172
Tunnel Services PIC....362

tunnel-limit statement.....190
usage guidelines.....172

U

unicast-umh-election statement.....190
upstream multicast hop
configuring.....175
upstream-interface statement.....272

V

version statement
BFD.....491
IGMP.....92
usage guidelines.....62
MLD.....135
usage guidelines.....101
PIM.....492
usage guidelines.....372, 378, 399
vlan statement
IGMP snooping.....313
vlan statement, IGMP snooping
usage guidelines.....298
VLANs
IGMP snooping.....298
VPN control plane.....166
vpn-group-address statement.....492
usage guidelines.....418

Index of Statements and Commands

A

accounting statement
- IGMP....75
- IGMP interface....75
- MLD....117
- MLD interface....117

active-source-limit statement....516
address statement
- anycast RPs....445
- local RPs....446
- static RPs....447

anycast-pim statement....449
asm-override-ssm statement....245
assert-timeout statement....450
authentication-key statement
- MSDP....517

auto-rp statement....452
autodiscovery statement....177
autodiscovery-only statement....178

B

backup-pe-group statement....246
backups statement....247
bandwidth statement....248
bfd-liveness-detection statement
- PIM....453

bootstrap statement....454
bootstrap-export statement....454
bootstrap-import statement....455
bootstrap-priority statement....456

D

data-encapsulation statement....518
data-mdt-reuse statement....179
default-peer statement....519
dense-groups statement....456
disable statement
- DVMRP....329
- IGMP....76
- MLD....118
- MSDP....520
- PIM family....458
- PIM graceful restart....457
- PIM interfaces....458
- PIM protocol....458
- SAP....143

dr-election-on-p2p statement....459
dr-register-policy statement....459
dvmrp statement....330

E

embedded-rp statement....460
exclude statement
- IGMP....76
- MLD....118

export statement
- DVMRP....331
- MSDP....521
- PIM....460, 461

F

family statement
- bootstrap....461
- local RP....463

flood-groups statement
- multicast snooping....283

flow-map statement....249
forwarding-cache statement
- multicast snooping....284

G

graceful-restart statement
- multicast snooping....285
- PIM....464

group statement
- IGMP....77
- IGMP snooping....299
- MDT....179
- MLD....119
- MSDP....522

group-address statement....180
group-count statement
- IGMP....77
- MLD....120

group-increment statement
IGMP....78
MLD....120
group-limit statement
IGMP snooping....300
group-policy statement
IGMP....78
MLD....121
group-range statement....180
group-ranges statement....465

H

hello-interval statement
PIM....466
hold-time statement
DVMRP....331
PIM....467
host-only-interface statement....301

I

igmp statement....79
igmp-snooping statement....302
ignore-stp-topology-change statement....285
immediate-leave statement
IGMP....80
IGMP snooping....303
MLD....122
import statement
bootstrap....467
DVMRP....332
MSDP....523
PIM....468
inclusive statement....181
inet-mdt statement
autodiscovery....181
infinity statement....468
interface statement
DVMRP....332
IGMP....81
IGMP snooping....304
MLD....123
multicast....252
multicast scoping....253
PIM....469
intra-as statement....182

J

join-load-balance statement....470
join-prune-timeout statement....470

L

listen statement....144
local statement
PIM....472
local-address statement....253
MSDP group....524
MSDP peer....524
PIM....473

M

mapping-agent-election statement....474
maximum statement....525
maximum-bandwidth statement....254
maximum-rps statement....474
maximum-transmit-rate statement
IGMP....82
MLD....124
mdt statement....183
metric statement
DVMRP....333
minimum-interval statement
PIM....475
minimum-receive-interval statement
PIM....453, 475
mld statement....125
mode statement
DVMRP....333
MSDP....526
PIM....476
msdp statement....527
multicast statement....255
multicast-router-interface statement....305
multicast-snooping-options statement....286
multiplier statement
PIM....453, 477
mvpn statement....184, 185

N

neighbor-policy statement
PIM....477
no-accounting statement
IGMP....75
MLD....117
no-qos-adjust statement....257

O

oif-map statement
IGMP (interface)....83
MLD (interface)....127

P

passive statement
IGMP (interface)....84
MLD (interface)....127

peer statement
MSDP....529
pgm statement....157
pim statement....478
pim-ssm statement....185
pim-to-igmp-proxy statement....258
pim-to-mld-proxy statement....259
prefix statement....261
priority statement
bootstrap....480
PIM....481
promiscuous-mode statement....85
proxy statement
IGMP snooping....306

Q

query-interval statement
IGMP....85
IGMP snooping....307
MLD....128
query-last-member-interval statement
IGMP....86
IGMP snooping....308
MLD....128
query-response-interval statement
IGMP....86
IGMP snooping....309
MLD....129

R

rate statement
MDT....186
redundant-sources statement....262
restart-duration statement....482
reverse-oif-mapping statement....263
rib-group statement
DVMRP....334
MSDP....530
PIM....482
robust-count statement
IGMP....87
IGMP snooping....310
MLD....129
rp statement....483
rp-register-policy statement....484
rp-set statement....485
rpf-check-policy statement....264

S

sap statement....145
scope statement....264
scope-policy statement....265
signaling statement....187
source statement
IGMP....87
IGMP snooping....310
MDT....188
MLD....130
MSDP....531
SSM....266
source-address statement
IGMP snooping....311
source-count statement
IGMP....88
MLD....130
source-increment statement
IGMP....88
MLD....131
spt-threshold statement....486
ssm-groups statement....266
ssm-map statement
IGMP....89
MLD....131
SSM....267
static statement
IGMP....89
IGMP snooping....312
MLD....132
PIM....487
subscriber-leave-timer statement....268

T

threshold statement
forwarding cache....269
MDT....189
MSDP....532
multicast snooping....287
timeout statement
flow map....270
forwarding cache....271
traceoptions statement
DVMRP....335
IGMP....90
MLD....133
MSDP....533
PGM....158
PIM....488
tunnel-limit statement....190

U

unicast-umh-election statement....190
upstream-interface statement....272

V

version statement
BFD.....491
IGMP.....92
MLD.....135
PIM.....492
vlan statement
IGMP snooping.....313
vpn-group-address statement.....492