



JUNOS® Software

Policy Framework Configuration Guide

Release 9.3

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089
USA

408-745-2000

www.juniper.net

Part Number: 530-027199-01 , Revision 1

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 Policy Framework Configuration Guide,
Release 9.3

Copyright © 2008, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Writing: Ines Salazar, Fran Singer, Alan Twigg, Lisa Kelly
Editing: Joanne McClintock, Nancy Kurahashi
Illustration: Faith Bradford, Nathaniel Woodward
Cover Design: Edmonds Design

Revision History
10 Oct 2008—Revision 1

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 présentes confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattache, soient rédigés 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	xxv
Part 1	Policy Framework	
Chapter 1	Policy Framework Overview	3
Part 2	Routing Policies	
Chapter 2	Routing Policy Framework Overview	15
Chapter 3	Routing Policy Configuration Statements	35
Chapter 4	Routing Policy Configuration	39
Chapter 5	Extended Match Conditions Configuration	91
Chapter 6	Extended Actions Configuration	129
Chapter 7	Summary of Routing Policy Configuration Statements	143
Part 3	Firewall Filters	
Chapter 8	Firewall Filter Overview	159
Chapter 9	Firewall Filter Configuration	163
Chapter 10	Policer Overview	231
Chapter 11	Policer Configuration	233
Chapter 12	Summary of Firewall Filter and Policer Configuration Statements	257
Part 4	Traffic Sampling and Forwarding	
Chapter 13	Traffic Sampling and Forwarding Overview	277
Chapter 14	Traffic Sampling and Forwarding Configuration	279
Chapter 15	Summary of Traffic Sampling and Forwarding Options Configuration Statements	335
Part 5	Indexes	
	Index	439
	Index of Statements and Commands	449

Table of Contents

	About This Guide	xxv
	Objectives	xxv
	Audience	xxv
	Supported Routing Platforms	xxvi
	Using the Indexes	xxvi
	Using the Examples in This Manual	xxvi
	Merging a Full Example	xxvii
	Merging a Snippet	xxvii
	Documentation Conventions	xxviii
	List of Technical Publications	xxx
	Documentation Feedback	xxxvii
	Requesting Technical Support	xxxvii
Part 1	Policy Framework	
Chapter 1	Policy Framework Overview	3
	Router Flows Affected by Policies	3
	Policy Architecture	6
	Control Points	6
	Policy Components	7
	Default Policies and Actions	7
	Configuration Tasks	8
	Policy Configuration Recommendations	8
	Comparison of Routing Policies and Firewall Filters	8
Part 2	Routing Policies	
Chapter 2	Routing Policy Framework Overview	15
	Importing and Exporting	16
	Routing Tables Affected by Routing Policies	18
	Default Routing Policies and Actions	19
	Creating Routing Policies	22

Configuring a Routing Policy	22
Match Conditions	23
Named Match Conditions	24
Actions	25
Terms	25
Routing Policy Application	25
Routing Protocols	26
Routing Policy Application to Routing Protocols	27
Forwarding Table	28
Evaluating a Routing Policy	28
How a Routing Policy Is Evaluated	28
How a Routing Policy Chain Is Evaluated	29
How a Routing Policy Expression Is Evaluated	30
How a Routing Policy Subroutine Is Evaluated	30
Routing Policy Tests	32
Supported Standards and Drafts	33

Chapter 3 Routing Policy Configuration Statements 35

Minimum Routing Policy Configuration	36
Minimum Routing Policy Chain Configuration	36
Minimum Subroutine Configuration	37

Chapter 4 Routing Policy Configuration 39

Defining Routing Policies	39
Routing Policy Name	40
Terms	40
Match Conditions	41
Actions	46
Flow Control Actions	48
Actions That Manipulate Route Characteristics	49
Trace Action	53
Final Action	54
Default Action	54
Example: Configure the Default Action for a Policy	55
Route List Actions	56
Examples: Defining Routing Policies	56
Defining a Routing Policy from BGP to IS-IS	56
Using Routing Policy to Set a Preference	56
Importing and Exporting Access and Access-Internal Routes in a Routing Policy	57
Applying Routing Policies	58
Applying Routing Policies to a Routing Protocol	58
Applying a Routing Policy	58
Applying a Routing Policy Chain	59

Applying Policy Expressions	59
Side Effects of Omitting the "from" Statement from an Export Policy	64
Applying Routing Policies to the Forwarding Table	66
Examples: Applying Routing Policies	66
Examples: Routing Policy Configuration	67
Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS	68
Example: Redistributing OSPF Routes into BGP	68
Example: Exporting Direct Routes Into IS-IS	69
Example: Exporting Internal IS-IS Level 1 Routes to Level 2	69
Example: Exporting IS-IS Level 2 Routes to Level 1	70
Example: Assigning Different Forwarding Next-Hop LSPs to Different Destination Prefixes	70
Example: Grouping Destination Prefixes	71
Example: Grouping Source Prefixes	72
Example: Grouping Source and Destination Prefixes in a Forwarding Class	73
Example: Accepting Routes with Specific Destination Prefixes	74
Example: Accepting Routes from BGP with a Specific Destination Prefix	74
Example: ISP Network Case Study	75
Requesting a Single Default Route on the Customer 1 Router	77
Requesting Specific Routes on the Customer 2 Router	78
Configuring a Peer Policy on ISP Router 3	80
Configuring Private and Exchange Peers on ISP Router 1 and 2	82
Configuring Locally Defined Static Routes on the Exchange Peer 2 Router	85
Configuring Outbound and Generated Routes on the Private Peer 2 Router	85
Configuring the Discard Interface	87
Testing Routing Policies	88
Example: Testing a Routing Policy	88

Chapter 5

Extended Match Conditions Configuration 91

Configuring AS Path Regular Expressions	91
Defining AS Path Regular Expressions	92
Null AS Path	95
Example: Null AS Path	96
How AS Path Regular Expressions Are Evaluated	96
Examples: Configuring AS Path Regular Expressions	96
Configuring Communities	97
Defining Communities	98
Configuring the Community Attribute	98
Configuring the Extended Communities Attribute	107
Inverting Community Matches	109
Configuring Link Bandwidth	109
How Communities Are Evaluated	109

Configuring Prefix Lists	110
Prefix List and Route List Differences	110
Defining Prefix Lists	111
How a Prefix List Is Evaluated	112
Configuring a Prefix List Filter	112
Example: Configuring a Prefix List	113
Configuring Route Lists	114
Defining Route Lists	114
How a Route List Is Evaluated	117
How Prefix Order Affects Route List Evaluation	118
Common Configuration Problem with the Longest-Match Lookup	118
Examples: Configuring Route Lists	119
Example: Rejecting Routes with Specific Destination Prefixes and Mask Lengths	119
Example: Rejecting Routes with a Mask Length Greater than Eight	119
Example: Rejecting Routes with Mask Length Between 26 and 29	120
Example: Rejecting Routes from Specific Hosts	120
Example: Accepting Routes with a Defined Set of Prefixes	121
Example: Rejecting Routes with a Defined Set of Prefixes	121
Example: Rejecting Routes with Prefixes Longer than 24 Bits	122
Example: Rejecting PIM Multicast Traffic Joins	122
Example: Rejecting PIM Traffic	122
Configuring Subroutines	123
Defining Subroutines	123
Termination Actions	124
Example: Configuring a Subroutine	126
Configuring the Condition Statement	127

Chapter 6

Extended Actions Configuration

129

Configuring the AS Path Prepend Action	129
Configuring the AS Path Expand Action	130
Configuring the Class-of-Service Action	130
Configuring the Damping Action	130
Configuring Flap Damping Parameters	131
Defining Damping Action	133
Enabling BGP Route Flap Damping	134
Disabling Damping by Prefix	134
Example: Disabling by Prefix	134
Example: Configuring BGP Flap Damping	134
Configuring the Load-Balance Per-Packet Action	136
Load Balancing Based on the MPLS Label Information	138
Load Balancing Based on Layer 2 MAC Information	140
Examples: Configuring Per-Packet Load Balancing	141

Chapter 7	Summary of Routing Policy Configuration Statements	143
	apply-path	143
	as-path	144
	as-path-group	144
	community	145
	condition	147
	damping	148
	export	149
	import	151
	policy-options	152
	policy-statement	153
	prefix-list	155
	prefix-list-filter	155
 Part 3	 Firewall Filters	
 Chapter 8	 Firewall Filter Overview	 159
	Firewall Filter Components	160
	Supported Standards	161
 Chapter 9	 Firewall Filter Configuration	 163
	Minimum Firewall Filter Configuration	164
	Configuring Firewall Filters	165
	Configuring the Family Address Type	166
	Configuring the Filter Name	166
	Configuring the Filter Terms	166
	Configuring a Filter Match Statement	167
	Configuring a Filter Action Statement	168
	Example: Configure a Filter Action Statement	171
	Example: Set the DSCP Bit to 0	173
	How Firewall Filters Are Evaluated	173
	Filter Match Conditions	174
	Specifying Numeric Range Filter Match Conditions	175
	Specifying Address Filter Match Conditions	184
	Specifying Bit-Field Filter Match Conditions	188
	Specifying Class-Based Filter Match Conditions	190
	Filtering Smaller Packets	191
	How Firewall Filters Test a Packet's Protocol	191
	Example: Do Not Test Packet Protocol	192
	Configuring a Filter Within a Filter	193
	Example: Configuring a Filter Within A Filter	194
	Examples: Defining Firewall Filters	194
	Example: Blocking Telnet and SSH Access	195
	Example: Blocking TFTP Access	195
	Example: Accepting DHCP Packets with Specific Addresses	196

Example: Defining a Policier for a Destination Class	196
Example: Counting IP Option Packets	197
Example: Accepting OSPF Packets from Certain Addresses	198
Example: Matching Packets Based on Two Unrelated Criteria	199
Example: Counting Both Accepted and Rejected Packets	199
Example: Blocking TCP Connections to a Certain Port Except from BGP Peers	200
Example: Accepting Packets with Specific IPv6 TCP Flags	201
Example: Setting a Rate Limit for Incoming Layer 2 Control Packets	201
Configuring Service Filters	203
Configuring Simple Filters	203
Example: Configuring a Simple Filter	204
Applying Firewall Filters to Interfaces	205
Configuring Interface-Specific Counters	206
Example: Configuring Interface-Specific Counters	206
Defining Interface Groups	207
Example: Defining Interface Groups	207
Configuring Firewall Filters for Logical Systems	208
Guidelines for Firewall Configuration in Logical Systems	209
Scenario 1: Firewall Objects That Reference Other Firewall Objects	210
Scenario 2: Referencing Firewall Objects from Outside the Firewall Configuration	211
Scenario 3: Firewall Objects That Reference Objects Outside the Firewall Configuration	215
Unsupported Configuration Statements, Actions, and Action Modifiers	216
Configuring Accounting	221
Configuring a Firewall Filter Accounting Profile	221
Configuring Filter-Based Forwarding	222
Examples: Configuring Filter-Based Forwarding	223
Configuring Forwarding Table Filters	225
Overview of Forwarding Table Filters	225
Configuring a Forwarding Table Filter	225
Configuring Firewall Filter System Logging	227
Example: Configuring Firewall Filter System Logging	228

Chapter 10	Policer Overview	231
-------------------	-------------------------	------------

Chapter 11	Policer Configuration	233
-------------------	------------------------------	------------

Minimum Policier Configuration	234
Configuring Policers	235
Configuring Rate Limiting	235
Configuring a Policier Action	236
Example: Configuring a Policier Action	237

Configuring Multifield Classification and Policing	237
Configuring Filter-Specific Policers	239
Configuring Prefix-Specific Actions	239
Examples: Configuring Prefix-Specific Actions	240
Examples: Classifying Traffic	244
Configuring an Interface Set	245
Applying an Interface Policer	245
Example: Applying an Interface Policer	246
Configuring an Aggregate Policer	246
Example: Configuring an Aggregate Policer	246
Configuring a Bandwidth Policer	247
Example: Configuring a Bandwidth Policer	247
Configuring a Load-Balance Group	248
Configuring Tricolor Marking	248
Configuring a Tricolor Marking Policer	249
Example: Configuring a Tricolor Marking Policer	250
Configuring an Interface Policer Using Tricolor Marking Policing	251
Example: Rate-Limiting Bandwidth Using Tricolor Marking Policing	251
Examples: Configuring Policing	252

Chapter 12

Summary of Firewall Filter and Policer Configuration Statements 257

accounting-profile	257
action	258
family	259
filter	260
filter-specific	261
firewall	261
if-exceeding	262
interface-set	263
interface-specific	263
load-balance-group	264
logical-bandwidth-policer	264
logical-interface-policer	265
policer	266
prefix-action	267
service-filter	268
simple-filter	269
term	270
three-color-policer	272
three-color-policer (Applying)	272
three-color-policer (Configuring)	273
virtual-channel	274

Part 4**Traffic Sampling and Forwarding****Chapter 13****Traffic Sampling and Forwarding Overview****277****Chapter 14****Traffic Sampling and Forwarding Configuration****279**

Minimum Traffic Sampling or Forwarding Configuration	285
Configuring a Forwarding Table Filter	286
Configuring IPv6 Accounting	287
Configuring Traffic Sampling	288
Configuring Discard Accounting	289
Configuring Flow Monitoring	290
Configuring a Next-Hop Group	291
Configuring Per-Flow Load-Balancing Information	292
Configuring Per-Prefix Load Balancing	293
Configuring Per-Flow Load Balancing Based on Hash Values	293
Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent	294
Configuring DNS and TFTP Packet Forwarding	295
Tracing BOOTP, DNS, and TFTP Forwarding Operations	296
Configuring the Log Filename	297
Configuring the Number and Size of Log Files	297
Configuring Access to the Log File	297
Configuring a Regular Expression for Lines to Be Logged	298
Configuring the Trace Operations	298
Example: Configuring DNS Packet Forwarding	299
Configuring the Extended DHCP Relay Agent	299
Interaction Between the DHCP Relay Agent, DHCP Client, and DHCP Servers	303
Access and Access-Internal Routes	303
DHCP State Persistence	304
Graceful Routing Engine Switchover	304
Overriding the Default DHCP Relay Configuration	304
Overwriting giaddr Information	305
Overriding Option 82 Information	305
Using Layer 2 Unicast Transmission for DHCP Packets	305
Trusting Option 82 Information	306
Disabling DHCP Relay	306
Using Option 60 Information to Forward Client Traffic to Specific DHCP Servers	306
Using Matching Option 60 Strings to Process DHCP Client Traffic	306
Using Nonmatching Option 60 Strings to Process DHCP Client Traffic	308
Displaying a Count of Discarded DHCP Packets with Option 60 Information	308
Enabling and Disabling Insertion of Option 82 Information	308
Configuring Agent-Circuit-Id Information	309
Configuring an Option 82 Prefix	309

Configuring Server Groups	311
Configuring Active Server Groups	311
Grouping Interfaces with Common DHCP Relay Configuration	312
Configuring Group-Specific DHCP Relay Options	312
Enabling the DHCP Relay Agent on Specified Interfaces	313
Using External AAA Authentication Services	313
Verifying and Managing DHCP Relay Agent Clients	314
Tracing Extended DHCP Relay Agent Operations	314
Configuring the Extended DHCP Relay Agent Log Filename	315
Configuring the Number and Size of Extended DHCP Relay Agent Log Files	315
Configuring Access to the Extended DHCP Relay Agent Log File	316
Configuring a Regular Expression for Extended DHCP Relay Agent Lines to Be Logged	316
Configuring the Extended DHCP Relay Agent Tracing Flags	316
Example: Minimum DHCP Relay Agent Configuration	317
Example: DHCP Relay Agent Configuration with Multiple Clients and Servers	318
Example: Using Option 60 Strings to Forward DHCP Client Traffic	319
Example: Using Option 60 Strings to Drop DHCP Client Traffic	320
Disabling Traffic Sampling	321
Examples: Configuring Traffic Sampling	321
Sampling a Single SONET/SDH Interface	321
Sampling All Traffic from a Single IP Address	322
Sampling All FTP Traffic	323
Configuring Traffic Sampling Output	324
Traffic Sampling Output Files	325
Tracing Traffic Sampling Operations	326
Configuring Flow Aggregation (cflowd)	326
Debugging cflowd Flow Aggregation	328
Configuring Active Flow Monitoring Using Version 9	329
Example: Configuring Active Flow Monitoring Using Version 9	330
Configuring Port Mirroring	330
Configuring Packet Capture	332

Chapter 15

Summary of Traffic Sampling and Forwarding Options Configuration Statements	335
accounting	336
active-server-group	337
aggregation	338
always-write-giaddr	339
always-write-option-82	340
authentication	341
autonomous-system-type	342
bootp	343
cflowd	344
cflowd (Discard Accounting)	344
cflowd (Flow Monitoring)	345
cflowd (Sampling)	346

circuit-id	347
circuit-type	348
client-response-ttl	348
default-local-server-group	349
default-relay-server-group	350
delimiter	351
description	352
description (Interface)	352
description (Service)	352
dhcp-relay	353
disable	356
disable (Packet Capture)	356
disable (Sampling)	356
disable-relay	357
domain	357
domain-name	358
drop	359
export-format	360
family	361
family (Filtering)	361
family (Sampling)	362
family inet	363
family inet (Load Balancing)	363
family inet (Monitoring)	364
family mpls	365
family multiservice	366
file	367
file (Helpers Trace Options)	367
file (Packet Capture)	368
file (Sampling)	368
file (Trace Options)	369
filename	370
filename (Packet Capture)	370
filename (Sampling)	370
files	371
files (Packet Capture)	371
files (Sampling)	371
filter	372
flood	372
flow-active-timeout	373
flow-export-destination	373
flow-inactive-timeout	374
forwarding-options	374
group	375
hash-key	377
helpers	379
indexed-next-hop	380
input	381
input (Forwarding Table)	381
input (Port Mirroring)	381
input (Sampling)	382

interface	383
interface (Accounting or Sampling)	383
interface (BOOTP)	384
interface (DHCP Relay Agent)	385
interface (DNS and TFTP Packet Forwarding or Relay Agent)	386
interface (Monitoring)	387
interface (Next-Hop Group)	387
interface (Port Mirroring)	388
layer2-unicast-replies	388
load-balance	389
local-dump	390
local-server-group	391
logical-system-name	392
mac-address	393
maximum-capture-size	393
maximum-hop-count	394
max-packets-per-second	394
minimum-wait-time	395
monitoring	396
next-hop	397
next-hop-group	397
no-filter-check	398
no-listen	398
no-local-dump	398
no-stamp	398
no-world-readable	399
option-60	399
option-82	400
output	401
output (Accounting)	401
output (Forwarding Table)	402
output (Monitoring)	402
output (Port Mirroring)	403
output (Sampling)	404
overrides	405
packet-capture	406
password	407
per-flow	408
per-prefix	409
port	409
port-mirroring	410
prefix	411
rate	412
relay-option-60	413
relay-option-82	414
relay-server-group	415
route-accounting	416
routing-instance	416
routing-instance-name	417
run-length	417
sampling	418

server	420
server (DHCP or BOOTP Service)	420
server (DNS and TFTP Service)	420
server-group	421
size	422
size (Packet Capture)	422
size (Sampling)	423
stamp	423
tftp	424
traceoptions	425
traceoptions (DNS and TFTP Packet Forwarding)	426
traceoptions (Extended DHCP Relay Agent)	428
traceoptions (Port Mirroring and Traffic Sampling)	430
trust-option-82	430
username-include	431
user-prefix	432
vendor-option	433
version	434
version9	435
world-readable	435

Part 5

Indexes

Index	439
Index of Statements and Commands	449

List of Figures

Figure 1: Flows of Routing Information and Packets	4
Figure 2: Routing Policies to Control Routing Information Flow	5
Figure 3: Firewall Filters to Control Packet Flow	5
Figure 4: Policy Control Points	6
Figure 5: Importing and Exporting Routes	17
Figure 6: Importing and Exporting Routing Policies	22
Figure 7: Routing Policy Components	23
Figure 8: Routing Policy Evaluation	29
Figure 9: Routing Policy Chain Evaluation	30
Figure 10: Routing Policy Subroutine Evaluation	32
Figure 11: ISP Network Example	76
Figure 12: Incoming and Outgoing Interface Policers	245
Figure 13: Configure Sampling Rate	288

List of Tables

Table 1: Notice Icons	xxviii
Table 2: Text and Syntax Conventions	xxviii
Table 3: Technical Documentation for Supported Routing Platforms	xxx
Table 4: JUNOS Software Network Operations Guides	xxxiv
Table 5: JUNOS Software with Enhanced Services Documentation	xxxv
Table 6: Additional Books Available Through http://www.juniper.net/books	xxxvi
Table 7: Purpose of Routing Policies and Firewall Filters	9
Table 8: Implementation Differences Between Routing Policies and Firewall Filters	9
Table 9: Protocols That Can Be Imported to and Exported from the Routing Table	18
Table 10: Routing Tables Affected by Routing Policies	18
Table 11: Default Routing Policies	19
Table 12: Match Conditions	23
Table 13: Apply Routing Policies to Protocols	27
Table 14: Routing Policy Match Conditions	42
Table 15: Flow Control Actions	48
Table 16: Actions That Manipulate Route Characteristics	49
Table 17: Policy Action Conversion Values	60
Table 18: Policy Expression Logical Operators	60
Table 19: AS Path Regular Expression Operators	94
Table 20: Examples of Defining AS Path Regular Expressions	94
Table 21: Community Attribute Regular Expression Operators	101
Table 22: Examples of Defining Community Attribute Regular Expressions	102
Table 23: Prefix List and Route List Differences	111
Table 24: Route List Match Types for a Prefix List Filter	113
Table 25: Route List Match Types for a Prefix List	115
Table 26: Match Type Examples	116
Table 27: Damping Parameters	132
Table 28: Firewall Filter Actions and Action Modifiers	169
Table 29: Numeric Range IPv4 Firewall Filter Match Conditions	176
Table 30: Numeric Range IPv6 Firewall Filter Match Conditions	179
Table 31: Numeric Range Layer 2 Bridging Filter Match Conditions (MX-series Only)	182
Table 32: Address Firewall Filter Match Conditions	186
Table 33: Bit-Field Firewall Filter Match Conditions	189
Table 34: Bit-Field Logical Operators	190
Table 35: Unsupported Firewall Statements for Logical Systems	217
Table 36: Unsupported Firewall Actions and Action Modifiers for Logical Systems	218

Table 37: Tracing Flags298

About This Guide

This preface provides the following guidelines for using the *JUNOS® Software Policy Framework Configuration Guide*:

- Objectives on page xxv
- Audience on page xxv
- Supported Routing Platforms on page xxvi
- Using the Indexes on page xxvi
- Using the Examples in This Manual on page xxvi
- Documentation Conventions on page xxviii
- List of Technical Publications on page xxx
- Documentation Feedback on page xxxvii
- Requesting Technical Support on page xxxvii

Objectives

This guide is designed for network administrators who are configuring and monitoring a Juniper Networks J-series, M-series, MX-series, or T-series routing platform.



NOTE: This guide documents Release 9.3 of the JUNOS software. For additional information about the 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 routing platform.

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 Routing Platforms

For the features described in this manual, the JUNOS software currently supports the following routing 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.xml; }
```

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

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

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

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

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

Documentation Conventions

Table 1 on page xxviii 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 xxviii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

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

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

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

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

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

List of Technical Publications

Table 3 on page xxx lists the software and hardware guides and release notes for Juniper Networks J-series, M-series, MX-series, and T-series routing platforms and describes the contents of each document. Table 4 on page xxxiv lists the books included in the *Network Operations Guide* series. Table 5 on page xxxv lists the manuals and release notes supporting JUNOS software with enhanced services. All documents are available at <http://www.juniper.net/techpubs/>.

Table 6 on page xxxvi lists additional books on Juniper Networks solutions that you can order through your bookstore. A complete list of such books is available at <http://www.juniper.net/books>.

Table 3: Technical Documentation for Supported Routing Platforms

Book	Description
JUNOS Software for Supported Routing Platforms	
<i>Access Privilege</i>	Explains how to configure access privileges in user classes by using permission flags and regular expressions. Lists the permission flags along with their associated command-line interface (CLI) operational mode commands and configuration statements.
<i>Class of Service</i>	Provides an overview of the class-of-service (CoS) functions of the JUNOS software and describes how to configure CoS features, including configuring multiple forwarding classes for transmitting packets, defining which packets are placed into each output queue, scheduling the transmission service level for each queue, and managing congestion through the random early detection (RED) algorithm.
<i>CLI User Guide</i>	Describes how to use the JUNOS command-line interface (CLI) to configure, monitor, and manage Juniper Networks routing platforms. This material was formerly covered in the <i>JUNOS System Basics Configuration Guide</i> .
<i>Feature Guide</i>	Provides a detailed explanation and configuration examples for several of the most complex features in the JUNOS software.
<i>High Availability</i>	Provides an overview of hardware and software resources that ensure a high level of continuous routing platform operation and describes how to configure high availability (HA) features such as nonstop active routing (NSR) and graceful Routing Engine switchover (GRES).
<i>MPLS Applications</i>	Provides an overview of traffic engineering concepts and describes how to configure traffic engineering protocols.

Table 3: Technical Documentation for Supported Routing Platforms (*continued*)

Book	Description
<i>Multicast Protocols</i>	Provides an overview of multicast concepts and describes how to configure multicast routing protocols.
<i>Multiplay Solutions</i>	Describes how you can deploy IPTV and voice over IP (VoIP) services in your network.
<i>MX-series Layer 2 Configuration Guide</i>	Provides an overview of the Layer 2 functions of the MX-series routers, including configuring bridging domains, MAC address and VLAN learning and forwarding, and spanning-tree protocols. It also details the routing instance types used by Layer 2 applications. All of this material was formerly covered in the <i>JUNOS Routing Protocols Configuration Guide</i> .
<i>MX-series Solutions Guide</i>	Describes common configuration scenarios for the features supported on the MX-series routers, including basic bridged VLANs with normalized VLAN tags, aggregated Ethernet links, bridge domains, Multiple Spanning Tree Protocol (MSTP), and integrated routing and bridging (IRB).
<i>Network Interfaces</i>	Provides an overview of the network interface functions of the JUNOS software and describes how to configure the network interfaces on the routing platform.
<i>Network Management</i>	Provides an overview of network management concepts and describes how to configure various network management features, such as SNMP and accounting options.
<i>Policy Framework</i>	Provides an overview of policy concepts and describes how to configure routing policy, firewall filters, and forwarding options.
<i>Protected System Domain</i>	Provides an overview of the JCS 1200 platform and the concept of Protected System Domains (PSDs). The JCS 1200 platform, which contains up to 12 Routing Engines running JUNOS software, can be connected to up to three T-series routing platforms. To configure a PSD, you assign any number of Flexible PIC concentrators (FPCs) on a T-series routing platform to a pair of Routing Engines on the JCS 1200 platform. Each PSD has the same capabilities and functionality as a physical router, with its own control plane, forwarding plane, and administration.
<i>Routing Protocols</i>	Provides an overview of routing concepts and describes how to configure routing, routing instances, and unicast routing protocols.
<i>Secure Configuration Guide for Common Criteria and JUNOS-FIPS</i>	Provides an overview of secure Common Criteria and JUNOS-FIPS protocols for the JUNOS software and describes how to install and configure secure Common Criteria and JUNOS-FIPS on a routing platform.
<i>Services Interfaces</i>	Provides an overview of the services interfaces functions of the JUNOS software and describes how to configure the services interfaces on the router.
<i>Software Installation and Upgrade Guide</i>	Describes the JUNOS software components and packaging and explains how to initially configure, reinstall, and upgrade the JUNOS system software. This material was formerly covered in the <i>JUNOS System Basics Configuration Guide</i> .

Table 3: Technical Documentation for Supported Routing Platforms (*continued*)

Book	Description
<i>Subscriber Access</i>	Provides an overview of the subscriber access features of the JUNOS software and describes how to configure subscriber access support on the router, including dynamic profiles, class of service, AAA, and access methods.
<i>System Basics</i>	Describes Juniper Networks routing platforms and explains how to configure basic system parameters, supported protocols and software processes, authentication, and a variety of utilities for managing your router on the network.
<i>VPNs</i>	Provides an overview and describes how to configure Layer 2 and Layer 3 virtual private networks (VPNs), virtual private LAN service (VPLS), and Layer 2 circuits. Provides configuration examples.
JUNOS References	
<i>Hierarchy and RFC Reference</i>	Describes the JUNOS configuration mode commands. Provides a hierarchy reference that displays each level of a configuration hierarchy, and includes all possible configuration statements that can be used at that level. This material was formerly covered in the <i>JUNOS System Basics Configuration Guide</i> .
<i>Interfaces Command Reference</i>	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot interfaces.
<i>Routing Protocols and Policies Command Reference</i>	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot routing policies and protocols, including firewall filters.
<i>System Basics and Services Command Reference</i>	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot system basics, including commands for real-time monitoring and route (or path) tracing, system software management, and chassis management. Also describes commands for monitoring and troubleshooting services such as class of service (CoS), IP Security (IPsec), stateful firewalls, flow collection, and flow monitoring.
<i>System Log Messages Reference</i>	Describes how to access and interpret system log messages generated by JUNOS software modules and provides a reference page for each message.
J-Web User Guide	
<i>J-Web Interface User Guide</i>	Describes how to use the J-Web graphical user interface (GUI) to configure, monitor, and manage Juniper Networks routing platforms.
JUNOS API and Scripting Documentation	
<i>JUNOScript API Guide</i>	Describes how to use the JUNOScript application programming interface (API) to monitor and configure Juniper Networks routing platforms.
<i>JUNOS XML API Configuration Reference</i>	Provides reference pages for the configuration tag elements in the JUNOS XML API.

Table 3: Technical Documentation for Supported Routing Platforms (*continued*)

Book	Description
<i>JUNOS XML API Operational Reference</i>	Provides reference pages for the operational tag elements in the JUNOS XML API.
<i>NETCONF API Guide</i>	Describes how to use the NETCONF API to monitor and configure Juniper Networks routing platforms.
<i>JUNOS Configuration and Diagnostic Automation Guide</i>	Describes how to use the commit script and self-diagnosis features of the JUNOS software. This guide explains how to enforce custom configuration rules defined in scripts, how to use commit script macros to provide simplified aliases for frequently used configuration statements, and how to configure diagnostic event policies.
Hardware Documentation	
<i>Hardware Guide</i>	Describes how to install, maintain, and troubleshoot routing platforms and components. Each platform has its own hardware guide.
<i>PIC Guide</i>	Describes the routing platform's Physical Interface Cards (PICs). Each platform has its own PIC guide.
<i>DPC Guide</i>	Describes the Dense Port Concentrators (DPCs) for all MX-series routers.
JUNOScope Documentation	
<i>JUNOScope Software User Guide</i>	Describes the JUNOScope software graphical user interface (GUI), how to install and administer the software, and how to use the software to manage routing platform configuration files and monitor routing platform operations.
Advanced Insight Solutions (AIS) Documentation	
<i>Advanced Insight Solutions Guide</i>	Describes the Advanced Insight Manager (AIM) application, which provides a gateway between JUNOS devices and Juniper Support Systems (JSS) for case management and intelligence updates. Explains how to run AI-Scripts on Juniper Networks devices.
J-series Routing Platform Documentation	
<i>Getting Started Guide</i>	Provides an overview, basic instructions, and specifications for J-series routing platforms. The guide explains how to prepare your site for installation, unpack and install the router and its components, install licenses, and establish basic connectivity. Use the <i>Getting Started Guide</i> for your router model.
<i>Basic LAN and WAN Access Configuration Guide</i>	Explains how to configure the interfaces on J-series Services Routers for basic IP routing with standard routing protocols, ISDN backup, and digital subscriber line (DSL) connections.
<i>Advanced WAN Access Configuration Guide</i>	Explains how to configure J-series Services Routers in virtual private networks (VPNs) and multicast networks, configure data link switching (DLSw) services, and apply routing techniques such as policies, stateless and stateful firewall filters, IP Security (IPsec) tunnels, and class-of-service (CoS) classification for safer, more efficient routing.

Table 3: Technical Documentation for Supported Routing Platforms (*continued*)

Book	Description
<i>Administration Guide</i>	Shows how to manage users and operations, monitor network performance, upgrade software, and diagnose common problems on J-series Services Routers.
Release Notes	
<i>JUNOS Release Notes</i>	Summarize new features and known problems for a particular software release, provide corrections and updates to published JUNOS, JUNOScript, and NETCONF manuals, provide information that might have been omitted from the manuals, and describe upgrade and downgrade procedures.
<i>Hardware Release Notes</i>	Describe the available documentation for the routing platform and summarize known problems with the hardware and accompanying software. Each platform has its own release notes.
<i>JUNOScope Release Notes</i>	Contain corrections and updates to the published JUNOScope manual, provide information that might have been omitted from the manual, and describe upgrade and downgrade procedures.
<i>AIS Release Notes</i>	Summarize AIS new features and guidelines, identify known and resolved problems, provide information that might have been omitted from the manuals, and provide initial setup, upgrade, and downgrade procedures.
<i>AIS AI-Scripts Release Notes</i>	Summarize AI-Scripts new features, identify known and resolved problems, provide information that might have been omitted from the manuals, and provide instructions for automatic and manual installation, including deleting and rolling back.
<i>J-series Services Router Release Notes</i>	Briefly describe Services Router features, identify known hardware problems, and provide upgrade and downgrade instructions.

Table 4: JUNOS Software Network Operations Guides

Book	Description
<i>Baseline</i>	Describes the most basic tasks for running a network using Juniper Networks products. Tasks include upgrading and reinstalling JUNOS software, gathering basic system management information, verifying your network topology, and searching log messages.
<i>Interfaces</i>	Describes tasks for monitoring interfaces. Tasks include using loopback testing and locating alarms.
<i>MPLS</i>	Describes tasks for configuring, monitoring, and troubleshooting an example MPLS network. Tasks include verifying the correct configuration of the MPLS and RSVP protocols, displaying the status and statistics of MPLS running on all routing platforms in the network, and using the layered MPLS troubleshooting model to investigate problems with an MPLS network.

Table 4: JUNOS Software Network Operations Guides (continued)

Book	Description
<i>MPLS Log Reference</i>	Describes MPLS status and error messages that appear in the output of the <code>show mpls lsp extensive</code> command. The guide also describes how and when to configure Constrained Shortest Path First (CSPF) and RSVP trace options, and how to examine a CSPF or RSVP failure in a sample network.
<i>MPLS Fast Reroute</i>	Describes operational information helpful in monitoring and troubleshooting an MPLS network configured with fast reroute (FRR) and load balancing.
<i>Hardware</i>	Describes tasks for monitoring M-series and T-series routing platforms.

To configure and operate a J-series Services Router running JUNOS software with enhanced services, you must also use the configuration statements and operational mode commands documented in JUNOS configuration guides and command references. To configure and operate a WX Integrated Services Module, you must also use WX documentation.

Table 5: JUNOS Software with Enhanced Services Documentation

Book	Description
All Platforms	
<i>JUNOS Software Interfaces and Routing Configuration Guide</i>	Explains how to configure J-series interfaces for basic IP routing with standard routing protocols, ISDN service, firewall filters (access control lists), and class-of-service (CoS) traffic classification.
<i>JUNOS Software Security Configuration Guide</i>	Explains how to configure and manage security services such as stateful firewall policies, IP Security (IPsec) virtual private networks (VPNs), firewall screens, Network Address Translation (NAT), Public Key Cryptography, and Application Layer Gateways (ALGs).
<i>JUNOS Software Administration Guide</i>	Shows how to monitor J-series devices and routing operations, firewall and security services, system alarms and events, and network performance. This guide also shows how to administer user authentication and access, upgrade software, and diagnose common problems.
<i>JUNOS Software CLI Reference</i>	Provides the complete JUNOS software with enhanced services configuration hierarchy and describes the configuration statements and operational mode commands not documented in the standard JUNOS manuals.
J-series Only	
<i>JUNOS Software with Enhanced Services Design and Implementation Guide</i>	Provides guidelines and examples for designing and implementing IPsec VPNs, firewalls, and routing on J-series Services Routers running JUNOS software with enhanced services.

Table 5: JUNOS Software with Enhanced Services Documentation (continued)

Book	Description
<i>JUNOS Software with Enhanced Services Quick Start</i>	Explains how to quickly set up a J-series Services Router. This document contains router declarations of conformity.
<i>JUNOS Software with Enhanced Services J-series Services Router Hardware Guide</i>	Provides an overview, basic instructions, and specifications for J-series Services Routers. This guide explains how to prepare a site, unpack and install the router, replace router hardware, and establish basic router connectivity. This guide contains hardware descriptions and specifications.
<i>JUNOS Software with Enhanced Services Migration Guide</i>	Provides instructions for migrating an SSG device running ScreenOS software or a J-series Services Router running the JUNOS software to JUNOS software with enhanced services.
<i>WXC Integrated Services Module Installation and Configuration Guide</i>	Explains how to install and initially configure a WXC Integrated Services Module in a J-series Services Router for application acceleration.
<i>JUNOS Software with Enhanced Services Release Notes</i>	Summarizes new features and known problems for a particular release of JUNOS software with enhanced services on J-series Services Routers, including J-Web interface features and problems. The release notes also contain corrections and updates to the manuals and software upgrade and downgrade instructions for JUNOS software with enhanced services.

Table 6: Additional Books Available Through <http://www.juniper.net/books>

Book	Description
<i>Interdomain Multicast Routing</i>	Provides background and in-depth analysis of multicast routing using Protocol Independent Multicast sparse mode (PIM SM) and Multicast Source Discovery Protocol (MSDP); details any-source and source-specific multicast delivery models; explores multiprotocol BGP (MBGP) and multicast IS-IS; explains Internet Gateway Management Protocol (IGMP) versions 1, 2, and 3; lists packet formats for IGMP, PIM, and MSDP; and provides a complete glossary of multicast terms.
<i>JUNOS Cookbook</i>	Provides detailed examples of common JUNOS software configuration tasks, such as basic router configuration and file management, security and access control, logging, routing policy, firewalls, routing protocols, MPLS, and VPNs.
<i>MPLS-Enabled Applications</i>	Provides an overview of Multiprotocol Label Switching (MPLS) applications (such as Layer 3 virtual private networks [VPNs], Layer 2 VPNs, virtual private LAN service [VPLS], and pseudowires), explains how to apply MPLS, examines the scaling requirements of equipment at different points in the network, and covers the following topics: point-to-multipoint label switched paths (LSPs), DiffServ-aware traffic engineering, class of service, interdomain traffic engineering, path computation, route target filtering, multicast support for Layer 3 VPNs, and management and troubleshooting of MPLS networks.
<i>OSPF and IS-IS: Choosing an IGP for Large-Scale Networks</i>	Explores the full range of characteristics and capabilities for the two major link-state routing protocols: Open Shortest Path First (OSPF) and IS-IS. Explains architecture, packet types, and addressing; demonstrates how to improve scalability; shows how to design large-scale networks for maximum security and reliability; details protocol extensions for MPLS-based traffic engineering, IPv6, and multipoint-to-multipoint routing; and covers troubleshooting for OSPF and IS-IS networks.

Table 6: Additional Books Available Through <http://www.juniper.net/books> (continued)

Book	Description
<i>Routing Policy and Protocols for Multivendor IP Networks</i>	Provides a brief history of the Internet, explains IP addressing and routing (Routing Information Protocol [RIP], OSPF, IS-IS, and Border Gateway Protocol [BGP]), explores ISP peering and routing policies, and displays configurations for both Juniper Networks and other vendors' routers.
<i>The Complete IS-IS Protocol</i>	Provides the insight and practical solutions necessary to understand the IS-IS protocol and how it works by using a multivendor, real-world approach.

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 name
- Document part number
- Page number
- Software release version (not required for *Network Operations Guides [NOGs]*)

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 located at <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

Policy Framework

- Policy Framework Overview on page 3

Chapter 1

Policy Framework Overview

The JUNOS software provides a *policy framework*, which is a collection of JUNOS policies that allows you to control flows of routing information and packets. The policy framework is composed of the following policies:

- Routing policy—Allows you to control the routing information between the routing protocols and the routing tables and between the routing tables and the forwarding table
- Firewall filter policy—Allows you to control packets transiting the router to a network destination and packets destined for and sent by the router



NOTE: The term *firewall filter policy* is used here to emphasize that a firewall filter is a policy and shares some fundamental similarities with a routing policy. However, when referring to a firewall filter policy in the rest of this manual, the term *firewall filter* is used.

This chapter discusses the following topics related to understanding the JUNOS policy framework:

- Router Flows Affected by Policies on page 3
- Policy Architecture on page 6
- Comparison of Routing Policies and Firewall Filters on page 8

Router Flows Affected by Policies

The JUNOS policies affect the following router flows:

- Flow of routing information between the routing protocols and the routing tables and between the routing tables and the forwarding table. The Routing Engine handles this flow. *Routing information* is the information about routes learned by the routing protocols from a router's neighbors. This information is stored in routing tables and is subsequently advertised by the routing protocols to the router's neighbors. Routing policies allow you to control the flow of this information.
- Flow of data packets in and out of the router's physical interfaces. The Packet Forwarding Engine handles this flow. *Data packets* are chunks of data that transit the router as they are being forwarded from a source to a destination. When a router receives a data packet on an interface, it determines where to forward

the packet by looking in the forwarding table for the best route to a destination. The router then forwards the data packet toward its destination through the appropriate interface. Firewall filters allow you to control the flow of these data packets.

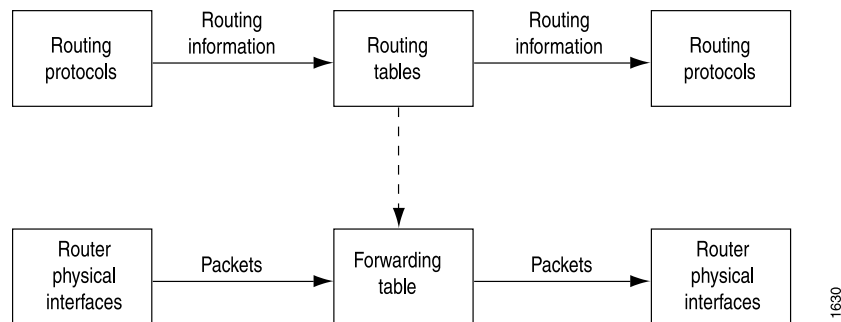
- Flow of local packets from the router's physical interfaces and to the Routing Engine. The Routing Engine handles this flow. *Local packets* are chunks of data that are destined for or sent by the router. Local packets usually contain routing protocol data, data for IP services such as telnet or SSH, and data for administrative protocols such as the Internet Control Message Protocol (ICMP). When the Routing Engine receives a local packet, it forwards the packet to the appropriate process or to the kernel, which are both part of the Routing Engine, or to the Packet Forwarding Engine. Firewall filters allow you to control the flow of these local packets.



NOTE: In the rest of this chapter, the term *packets* refers to both data and local packets unless explicitly stated otherwise.

Figure 1 on page 4 illustrates the flows through the router. Although the flows are very different from each other, they are also interdependent. Routing policies determine which routes are placed in the forwarding table. The forwarding table, in turn, has an integral role in determining the appropriate physical interface through which to forward a packet.

Figure 1: Flows of Routing Information and Packets




You can configure routing policies to control which routes the routing protocols place in the routing tables and to control which routes the routing protocols advertise from the routing tables (see Figure 2 on page 5). The routing protocols advertise active routes only from the routing tables. (An *active route* is a route that is chosen from all routes in the routing table to reach a destination. For information about the active route selection process, see the *JUNOS Routing Protocols Configuration Guide*.)

You can also use routing policies to do the following:

- Change specific route characteristics, which allow you to control which route is selected as the active route to reach a destination. In general, the active route is also advertised to a router's neighbors.
- Change to the default Border Gateway Protocol (BGP) route flap-damping values.

- Perform per-packet load balancing.
- Enable class of service (CoS).

Figure 2: Routing Policies to Control Routing Information Flow

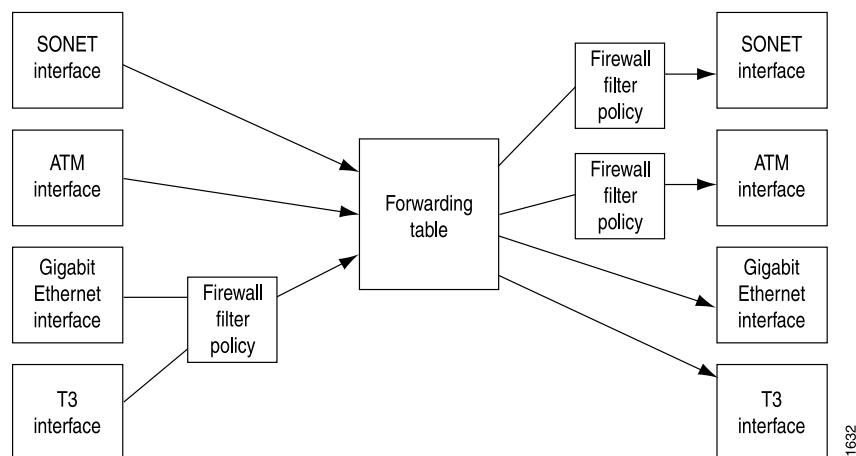


You can configure firewall filters to control the following (see Figure 3 on page 5):

- Which data packets are accepted on and transmitted from the physical interfaces. To control the flow of data packets, you apply firewall filters to the physical interfaces.
- Which local packets are transmitted from the physical interfaces and to the Routing Engine. To control local packets, you apply firewall filters on the loopback interface, which is the interface to the Routing Engine.

Firewall filters provide a means of protecting your router from excessive traffic transiting the router to a network destination or destined for the Routing Engine. Firewall filters that control local packets can also protect your router from external incidents such as denial-of-service attacks.

Figure 3: Firewall Filters to Control Packet Flow



Policy Architecture

A *policy* is a mechanism in the JUNOS policy framework that allows you to configure criteria against which something can be compared and an action that is performed if the criteria are met.

All policies in the JUNOS policy framework share the following architecture and configuration fundamentals:

- Control Points on page 6
- Policy Components on page 7
- Default Policies and Actions on page 7
- Configuration Tasks on page 8
- Policy Configuration Recommendations on page 8



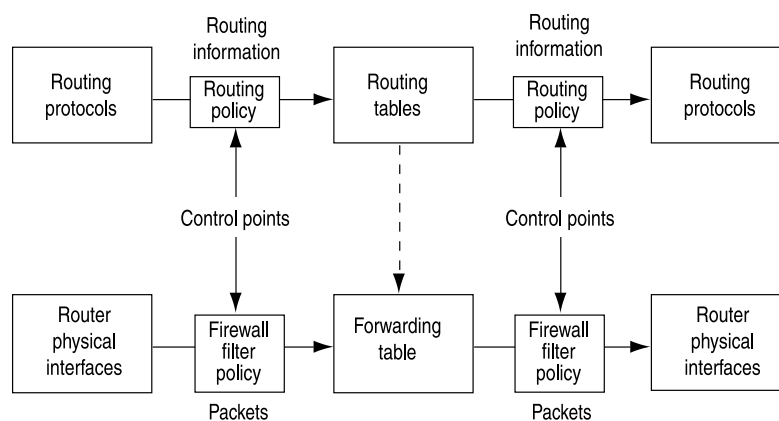
NOTE: This section highlights the fundamental architecture that all policies share. Note, however, that the implementation details of routing policies and firewall filters are very different. For information about these differences, see “Comparison of Routing Policies and Firewall Filters” on page 8.

Control Points

All policies provide two points at which you can control routing information or packets through the router (see Figure 4 on page 6). These control points allow you to control the following:

- Routing information before and after it is placed in the routing table.
- Data packets before and after a forwarding table lookup.
- Local packets before and after they are received by the Routing Engine. (Figure 4 on page 6 appears to depict only one control point but because of the bidirectional flow of the local packets, two control points actually exist.)

Figure 4: Policy Control Points



Because there are two control points, you can configure policies that control the routing information or data packets before and after their interaction with their respective tables, and policies that control local packets before and after their interaction with the Routing Engine. *Import routing policies* control the routing information that is placed in the routing tables, whereas *export routing policies* control the routing information that is advertised from the routing tables. *Input firewall filters* control packets that are received on a router interface, whereas *output firewall filters* control packets that are transmitted from a router interface.

Policy Components

All policies are composed of the following components that you configure:

- *Match conditions*—Criteria against which a route or packets are compared. You can configure one or more criteria. If all criteria match, one or more actions are applied.
- *Actions*—What happens if all criteria match. You can configure one or more actions.
- *Terms*—Named structures in which match conditions and actions are defined. You can define one or more terms.

For more information about these concepts and how they fit into the context of their respective policies, see “Configuring a Routing Policy” on page 22 and “Firewall Filter Components” on page 160.

The policy framework software evaluates each incoming and outgoing route or packet against the match conditions in a term. If the criteria in the match conditions are met, the defined action is taken.

In general, the policy framework software compares the route or packet against the match conditions in the first term in the policy, then goes on to the next term, and so on. (For specific information about when the evaluation process ends for each policy, see “Comparison of Routing Policies and Firewall Filters” on page 8.) Therefore, the order in which you arrange terms in a policy is relevant.

The order of match conditions within a term is not relevant because a route or packet must match all match conditions in a term for an action to be taken.

Default Policies and Actions

If an incoming or outgoing route or packet arrives and there is no explicitly configured policy related to the route or to the interface upon which the packet arrives, the action specified by the default policy is taken. A *default policy* is a rule or a set of rules that determine whether the route is placed in or advertised from the routing table, or whether the packet is accepted into or transmitted from the router interface.

All policies also have default actions in case one of the following situations arises during policy evaluation:

- A policy does not specify a match condition.
- A match occurs, but a policy does not specify an action.

- A match does not occur with a term in a policy and subsequent terms in the same policy exist.
- A match does not occur by the end of a policy.

Configuration Tasks

All policies share a two-step configuration process:

- Define the policy—Define the policy components. The components include criteria against which routes or packets are compared and actions that are performed if the criteria are met. For more information, see “Policy Components” on page 7.
- Apply the policy—Apply the policy to whatever moves the routing information or packets through the router, for example, the routing protocol or the router interface.



NOTE: A defined policy does not take effect until you apply it.

Policy Configuration Recommendations

The JUNOS policy architecture is simple and straightforward. However, the actual implementation of each policy adds layers of complexity to the policy as well as adding power and flexibility to your router’s capabilities. Configuring a policy has a major impact on the flow of routing information or packets within and through the router. For example, you can configure a routing policy that does not allow routes associated with a particular customer to be placed in the routing table. As a result of this routing policy, the customer routes are not used to forward data packets to various destinations and the routes are not advertised by the routing protocol to neighbors.

Before configuring a policy, determine what you want to accomplish with it and thoroughly understand how to achieve your goal using the various match conditions and actions. Also, make certain that you understand the default policies and actions for the policy you are configuring.

Comparison of Routing Policies and Firewall Filters

Although routing policies and firewall filters share an architecture, as described in “Policy Architecture” on page 6, their purposes, implementation, and configuration are different. Table 7 on page 9 describes their purposes. Table 8 on page 9 compares the implementation details for routing policies and firewall filters, highlighting the similarities and differences in their configuration.

For complete information about routing policies, see “Routing Policies” on page 13. For complete information about firewall filters, see “Firewall Filters” on page 157.

Table 7: Purpose of Routing Policies and Firewall Filters

Policies	Source	Policy Purpose
Routing policies	Routing information is generated by internal networking peers.	To control the size and content of the routing tables, which routes are advertised, and which routes are considered the best to reach various destinations.
Firewall filters	Packets are generated by internal and external devices through which hostile attacks can be perpetrated.	To protect your router and network from excessive incoming traffic or hostile attacks that can disrupt network service, and to control which packets are forwarded from which router interfaces.

Table 8: Implementation Differences Between Routing Policies and Firewall Filters

Policy Architecture	Routing Policy Implementation	Firewall Filter Implementation
Control points	Control routing information that is placed in the routing table with an import routing policy and advertised from the routing table with an export routing policy.	Control packets that are accepted on a router interface with an input firewall filter and that are forwarded from an interface with an output firewall filter.
Configuration tasks:	Define a policy that contains terms, match conditions, and actions.	Define a policy that contains terms, match conditions, and actions.
<ul style="list-style-type: none"> ■ Define policy ■ Apply policy 	<p>Apply one or more export or import policies to a routing protocol. You can also apply a <i>policy expression</i>, which uses Boolean logical operators with multiple import or export policies.</p> <p>You can also apply one or more export policies to the forwarding table.</p>	<p>Apply one input or output firewall filter to a physical interface or physical interface group to filter data packets received by or forwarded to a physical interface (on routing platforms with an Internet Processor II application-specific integrated circuit [ASIC] only).</p> <p>You can also apply one input or output firewall filter to the routing platform's loopback interface, which is the interface to the Routing Engine (on all routing platforms). This allows you to filter local packets received by or forwarded from the Routing Engine.</p>
Terms	<p>Configure as many terms as desired. Define a name for each term.</p> <p>Terms are evaluated in the order in which you specify them.</p> <p>Evaluation of a policy ends after a packet matches the criteria in a term and the defined or default policy action of accept or reject is taken. The route is not evaluated against subsequent terms in the same policy or subsequent policies.</p>	<p>Configure as many terms as desired. Define a name for each term.</p> <p>Terms are evaluated in the order in which you specify them.</p> <p>Evaluation of a firewall filter ends after a packet matches the criteria in a term and the defined or default action is taken. The packet is not evaluated against subsequent terms in the firewall filter.</p>

Table 8: Implementation Differences Between Routing Policies and Firewall Filters *(continued)*

Policy Architecture	Routing Policy Implementation	Firewall Filter Implementation
Match conditions	<p>Specify zero or more criteria that a route must match. You can specify criteria based on source, destination, or properties of a route. You can also specify the following match conditions, which require more configuration:</p> <ul style="list-style-type: none"> ■ Autonomous system (AS) path expression—A combination of AS numbers and regular expression operators. ■ Community—A group of destinations that share a common property. ■ Prefix list—A named list of prefixes. ■ Route list—A list of destination prefixes. ■ Subroutine—A routing policy that is called repeatedly from other routing policies. 	<p>Specify zero or more criteria that a packet must match. You must match various fields in the packet's header. The fields are grouped into the following categories:</p> <ul style="list-style-type: none"> ■ Numeric values, such as port and protocol numbers. ■ Prefix values, such as IP source and destination prefixes. ■ Bit-field values—Whether particular bits in the fields are set, such as IP options, Transmission Control Protocol (TCP) flags, and IP fragmentation fields. You can specify the fields using Boolean logical operators.
Actions	<p>Specify zero or one action to take if a route matches all criteria. You can specify the following actions:</p> <ul style="list-style-type: none"> ■ Accept—Accept the route into the routing table, and propagate it. After this action is taken, the evaluation of subsequent terms and policies ends. ■ Reject—Do not accept the route into the routing table, and do not propagate it. After this action is taken, the evaluation of subsequent terms and policies ends. <p>In addition to the preceding actions, you can also specify zero or more of the following types of actions:</p> <ul style="list-style-type: none"> ■ Next term—Evaluate the next term in the routing policy. ■ Next policy—Evaluate the next routing policy. ■ Actions that manipulate characteristics associated with a route as the routing protocol places it in the routing table or advertises it from the routing table. ■ Trace action, which logs route matches. 	<p>Specify zero or one action to take if a packet matches all criteria. (We recommend that you always explicitly configure an action.) You can specify the following actions:</p> <ul style="list-style-type: none"> ■ Accept—Accept a packet. ■ Discard—Discard a packet silently, without sending an ICMP message. ■ Reject—Discard a packet, and send an ICMP destination unreachable message. ■ Routing instance—Specify a routing table to which packets are forwarded. ■ Next term—Evaluate the next term in the firewall filter. <p>In addition to zero or the preceding actions, you can also specify zero or more action modifiers. You can specify the following action modifiers:</p> <ul style="list-style-type: none"> ■ Count—Add packet to a count total. ■ Forwarding class—Set the packet forwarding class to a specified value from 0 through 3. ■ IPSec security association—Used with the source and destination address match conditions, specify an IP Security (IPSec) security association (SA) for the packet. ■ Log—Store the header information of a packet on the Routing Engine. ■ Loss priority—Set the packet loss priority (PLP) bit to a specified value, 0 or 1. ■ Policer—Apply rate-limiting procedures to the traffic. ■ Sample—Sample the packet traffic. ■ Syslog—Log an alert for the packet.

Table 8: Implementation Differences Between Routing Policies and Firewall Filters *(continued)*

Policy Architecture	Routing Policy Implementation	Firewall Filter Implementation
Default policies and actions	<p>If an incoming or outgoing route arrives and a policy related to the route is not explicitly configured, the action specified by the default policy for the associated routing protocol is taken.</p> <p>The following default actions exist for routing policies:</p> <ul style="list-style-type: none"> ■ If a policy does not specify a match condition, all routes evaluated against the policy match. ■ If a match occurs but the policy does not specify an accept, reject, next term, or next policy action, one of the following occurs: <ul style="list-style-type: none"> ■ The next term, if present, is evaluated. ■ If no other terms are present, the next policy is evaluated. ■ If no other policies are present, the action specified by the default policy is taken. ■ If a match does not occur with a term in a policy and subsequent terms in the same policy exist, the next term is evaluated. ■ If a match does not occur with any terms in a policy and subsequent policies exist, the next policy is evaluated. ■ If a match does not occur by the end of a policy and no other policies exist, the accept or reject action specified by the default policy is taken. 	<p>If an incoming or outgoing packet arrives on an interface and a firewall filter is not configured for the interface, the default policy is taken (the packet is accepted).</p> <p>The following default actions exist for firewall filters:</p> <ul style="list-style-type: none"> ■ If a firewall filter does not specify a match condition, all packets are considered to match. ■ If a match occurs but the firewall filter does not specify an action, the packet is accepted. ■ If a match occurs, the defined or default action is taken and the evaluation ends. Subsequent terms in the firewall filter are not evaluated, unless the next term action is specified. ■ If a match does not occur with a term in a firewall filter and subsequent terms in the same filter exist, the next term is evaluated. ■ If a match does not occur by the end of a firewall filter, the packet is discarded.

Part 2

Routing Policies

- Routing Policy Framework Overview on page 15
- Routing Policy Configuration Statements on page 35
- Routing Policy Configuration on page 39
- Extended Match Conditions Configuration on page 91
- Extended Actions Configuration on page 129
- Summary of Routing Policy Configuration Statements on page 143

Chapter 2

Routing Policy Framework Overview

All routing protocols store their routing information in routing tables. From these tables, the routing protocols calculate the best route to each destination and place these routes in a forwarding table. These routes are then used to forward routing protocol traffic toward a destination, and they can be advertised to neighbors using one or more routing protocols.



NOTE: Instead of referring to the multiple routing tables that the JUNOS software maintains, the discussion in the rest of this chapter assumes the `inet.0` routing table unless explicitly stated otherwise. By default, the JUNOS software stores unicast IP version 4 (IPv4) routes in the `inet.0` routing table. For information about all the routing tables, see “Routing Tables Affected by Routing Policies” on page 18.

In general, the routing protocols place all their routes in the routing table and advertise a limited set of routes from the routing table. The general rules for handling the routing information between the routing protocols and the routing table are known as the *routing policy framework*.

The routing policy framework is composed of default rules for each routing protocol that determine which routes the protocol places in the routing table and advertises from the routing table. The default rules for each routing protocol are known as *default routing policies*.

You can create routing policies to preempt the default policies, which are always present. A *routing policy* is a mechanism in the JUNOS software that allows you to modify the routing policy framework to suit your needs. You can create and implement your own routing policies to do the following:

- Control which routes a routing protocol places in the routing table.
- Control which active routes a routing protocol advertises from the routing table. (An *active route* is a route that is chosen from all routes in the routing table to reach a destination. For information about the active route selection process, see the *JUNOS Routing Protocols Configuration Guide*.)
- Manipulate the route characteristics as a routing protocol places it in the routing table or advertises it from the routing table.

You can manipulate the route characteristics to control which route is selected as the active route to reach a destination. The active route is placed in the forwarding

table and used to forward traffic toward the route's destination. In general, the active route is also advertised to a router's neighbors.

To create a routing policy, you must define the policy and apply it. You define the policy by specifying the criteria that a route must match and the actions to perform if a match occurs. You then apply the policy to a routing protocol or to the forwarding table.



NOTE: Before you create your routing policies, we recommend that you read through this entire section to become familiar with the terminology, concepts, and configuration guidelines.

This chapter discusses the following topics related to understanding and creating routing policies:

- Importing and Exporting on page 16
- Default Routing Policies and Actions on page 19
- Creating Routing Policies on page 22
- Configuring a Routing Policy on page 22
- Evaluating a Routing Policy on page 28
- Routing Policy Tests on page 32
- Supported Standards and Drafts on page 33

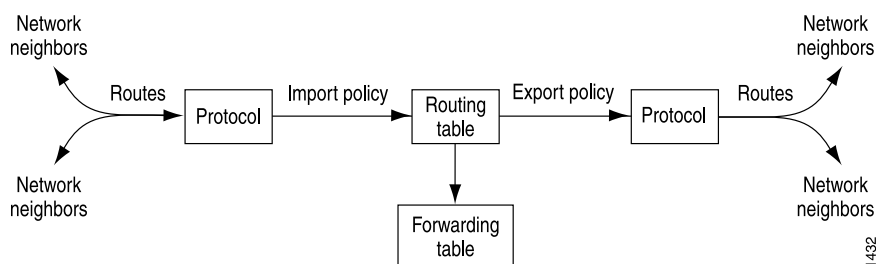
Importing and Exporting

Two terms—*import* and *export*—explain how routes move between the routing protocols and the routing table (see Figure 5 on page 17):

- When the Routing Engine places the routes of a routing protocol into the routing table, it is *importing* routes into the routing table.
- When the Routing Engine uses active routes from the routing table to send a protocol advertisement, it is *exporting* routes from the routing table.



NOTE: The process of moving routes between a routing protocol and the routing table is described always *from the point of view of the routing table*. That is, routes are *imported into* a routing table from a routing protocol and they are *exported from* a routing table to a routing protocol. Remember this distinction when working with routing policies.

Figure 5: Importing and Exporting Routes

When evaluating routes for export, the Routing Engine uses only active routes from the routing table. For example, if a routing table contains multiple routes to the same destination and one route has a preferable metric, only that route is evaluated. In other words, an export policy does not evaluate all routes; it evaluates only those routes that a routing protocol is allowed to advertise to a neighbor. For more information about the active path selection algorithm, see the *JUNOS Routing Protocols Configuration Guide*.



NOTE: By default, the Border Gateway Protocol (BGP) advertises active routes. However, you can configure BGP to advertise *inactive routes*, which go to the same destination as other routes but have less preferable metrics. For more information about advertising inactive routes, see the *JUNOS Routing Protocols Configuration Guide*.

Table 9 on page 18 lists the routing protocols from which the routing table can import routes and to which the routing table can export routes. Table 9 on page 18 also lists direct and explicitly configured routes, which for the purposes of this table are considered a pseudoprotocol. (An *explicitly configured route* is a route that you have configured. *Direct routes* are not explicitly configured; they are created as a result of IP addresses being configured on an interface.) Explicitly configured routes include aggregate, generated, local, and static routes. (An *aggregate route* is a route that distills groups of routes with common addresses into one route. A *generated route* is a route used when the routing table has no information about how to reach a particular destination. A *local route* is an IP address assigned to a router interface. A *static route* is an unchanging route to a destination.)

The policy framework software treats direct and explicitly configured routes as if they are learned through routing protocols; therefore, they can be imported into the routing table. Routes cannot be exported from the routing table to the pseudoprotocol, because this protocol is not a real routing protocol. However, aggregate, direct, generated, and static routes can be exported from the routing table to routing protocols, whereas local routes cannot.

For information about the default routing policies for each routing protocol, see Table 11 on page 19. For information about the import and export routing policies supported for each routing protocol and the level at which you can apply these policies, see Table 9 on page 18.

Table 9: Protocols That Can Be Imported to and Exported from the Routing Table

Protocol	Import	Export
BGP	Yes	Yes
Distance Vector Multicast Routing Protocol (DVMRP)	Yes	Yes
Intermediate System-to-Intermediate System (IS-IS)	Yes	Yes
Label Distribution Protocol (LDP)	Yes	Yes
Multiprotocol Label Switching (MPLS)	Yes	No
Open Shortest Path First (OSPF)	Yes	Yes
Protocol Independent Multicast (PIM) dense mode	Yes	Yes
PIM sparse mode	Yes	Yes
PIM sparse-dense mode	Yes	Yes
Pseudoprotocol:	Yes	No
<ul style="list-style-type: none"> ■ Direct routes ■ Explicitly configured routes <ul style="list-style-type: none"> ■ Aggregate routes ■ Generated routes ■ Local routes ■ Static routes 		
Routing Information Protocol (RIP) and Routing Information Protocol next generation (RIPng)	Yes	Yes

Routing Tables Affected by Routing Policies

Table 10 on page 18 lists the routing tables affected by default and user-defined routing policies and the types of routes that each routing table stores.

Table 10: Routing Tables Affected by Routing Policies

Routing Table	Type of Routes Stored
inet.0	Unicast IPv4 routes
.inet.0 <i>instance-name</i>	Unicast IPv4 routes for a particular routing instance
inet.1	Multicast IPv4 routes
inet.2	Unicast IPv4 routes for multicast reverse-path forwarding (RPF) lookup
inet.3	MPLS routes

Table 10: Routing Tables Affected by Routing Policies (*continued*)

Routing Table	Type of Routes Stored
mpls.0	MPLS routes for label-switched path (LSP) next hops
inet6.0	Unicast IP version 6 (IPv6) routes



NOTE: The discussion in the rest of this chapter assumes that the routing table is `inet.0` unless explicitly stated otherwise.

For more information about routing tables, see the *JUNOS Routing Protocols Configuration Guide*.

Default Routing Policies and Actions

You must be familiar with the default routing policies to know when you need to modify them to suit your needs. Table 11 on page 19 summarizes the for each routing protocol that imports and exports routes. The actions in the are taken if you have not explicitly configured a routing policy. This table also shows direct and explicitly configured routes, which for the purposes of this table are considered a pseudoprotocol. Explicitly configured routes include aggregate, generated, and static routes.

The default import policy is always the same: accept all routes learned from the protocol. Table 11 on page 19 includes information about the routing tables used by each protocol.

Table 11: Default Routing Policies

Importing or Exporting Protocol	Default Import Policy	Default Export Policy
BGP	Accept all BGP IPv4 routes learned from configured neighbors and import into the <code>inet.0</code> routing table. Accept all BGP IPv6 routes learned from configured neighbors and import into the <code>inet6.0</code> routing table.	Accept and export active BGP routes.
DVMRP	Accept all DVMRP routes and import into the <code>inet.1</code> routing table.	Accept and export active DVMRP routes.
IS-IS	Accept all IS-IS routes and import into the <code>inet.0</code> and <code>inet6.0</code> routing tables. (You cannot override or change this default policy.)	Reject everything. (The protocol uses flooding to announce local routes and any learned routes.)

Table 11: Default Routing Policies *(continued)*

Importing or Exporting Protocol	Default Import Policy	Default Export Policy
LDP	Accept all LDP routes and import into the <code>inet.3</code> routing table.	Reject everything.
MPLS	Accept all MPLS routes and import into the <code>inet.3</code> routing table.	Accept and export active MPLS routes.
OSPF	Accept all OSPF routes and import into the <code>inet.0</code> routing table. (You cannot override or change this default policy.)	Reject everything. (The protocol uses flooding to announce local routes and any learned routes.)
PIM dense mode	Accept all PIM dense mode routes and import into the <code>inet.1</code> routing table.	Accept active PIM dense mode routes.
PIM sparse mode	Accept all PIM sparse mode routes and import into the <code>inet.1</code> routing table.	Accept and export active PIM sparse mode routes.
Pseudoprotocol: ■ Direct routes ■ Explicitly configured routes: ■ Aggregate routes ■ Generated routes ■ Static routes	Accept all direct and explicitly configured routes and import into the <code>inet.0</code> routing table.	The pseudoprotocol cannot export any routes from the routing table because it is not a routing protocol. Routing protocols can export these or any routes from the routing table.
RIP	Accept all RIP routes learned from configured neighbors and import into the <code>inet.0</code> routing table.	Reject everything. To export RIP routes, you must configure an export policy for RIP.
RIPng	Accept all RIPng routes learned from configured neighbors and import into the <code>inet6.0</code> routing table.	Reject everything. To export RIPng routes, you must configure an export policy for RIPng.
Test policy	Accept all routes. For additional information about test policy, see “Routing Policy Tests” on page 32.	

When multiple routes for a destination exist in the routing table, the protocol selects an active route and that route is placed in the appropriate routing table. For equal-cost routes, the JUNOS software places multiple next hops in the appropriate routing table.

When a protocol is exporting routes from the routing table, it exports active routes only. This applies to actions specified by both default and user-defined export policies.

You cannot change the default import policy for the link-state protocols IS-IS and OSPF. As link-state protocols, IS-IS and OSPF exchange routes between systems within an autonomous system (AS). All routers and systems within an AS must share the same link-state database, which includes routes to reachable prefixes and the metrics associated with the prefixes. If an import policy is configured and applied to IS-IS or OSPF, some routes might not be learned or advertised or the metrics for learned routes might be altered, which would make a consistent link-state database impossible.

The default export policy for IS-IS and OSPF protocols is to reject everything. These protocols do not actually export their internally learned routes (the directly connected routes on interfaces that are running the protocol). Both IS-IS and OSPF protocols use a procedure called flooding to announce local routes and any routes learned by the protocol. The flooding procedure is internal to the protocol, and is unaffected by the policy framework. Exporting can be used only to announce information from other protocols, and the default is not to do so.

For information about the routing protocols from which the routing table can import routes and to which routing protocols the routing table can export routes, see Table 9 on page 18. For information about the user-defined import and export policies supported for each routing protocol and the level at which you can apply these policies, see Table 13 on page 27.

The following default actions are taken if the following situations arise during policy evaluation:

- If a policy does not specify a match condition, all routes evaluated against the policy match.
- If a match occurs but the policy does not specify an accept, reject, next term, or next policy action, one of the following occurs:
 - The next term, if present, is evaluated.
 - If no other terms are present, the next policy is evaluated.
 - If no other policies are present, the action specified by the default policy is taken.
- If a match does not occur with a term in a policy and subsequent terms in the same policy exist, the next term is evaluated.
- If a match does not occur with any terms in a policy and subsequent policies exist, the next policy is evaluated.
- If a match does not occur by the end of a policy or all policies, the accept or reject action specified by the default policy is taken.

Creating Routing Policies

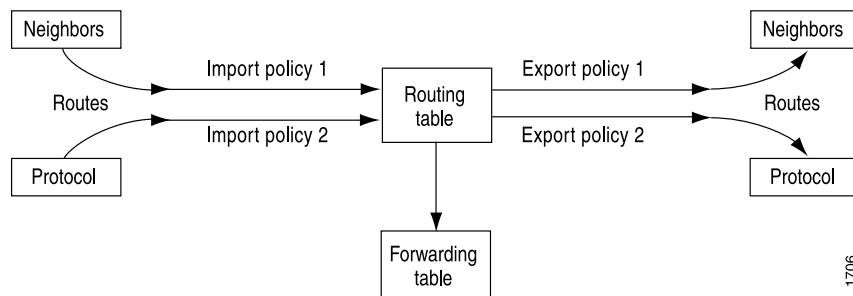
The following are typical circumstances under which you might want to preempt the default routing policies in the routing policy framework by creating your own routing policies:

- You do not want a protocol to import all routes into the routing table. If the routing table does not learn about certain routes, they can never be used to forward packets and they can never be redistributed into other routing protocols.
- You do not want a routing protocol to export all the active routes it learns.
- You want a routing protocol to announce active routes learned from another routing protocol, which is sometimes called *route redistribution*.
- You want to manipulate route characteristics, such as the preference value, AS path, or community. You can manipulate the route characteristics to control which route is selected as the active route to reach a destination. In general, the active route is also advertised to a router's neighbors.
- You want to change the default BGP route flap-damping parameters.
- You want to perform per-packet load balancing.
- You want to enable class of service (CoS).

Configuring a Routing Policy

As shown in Figure 6 on page 22, you use *import routing policies* to control which routes routing protocols place in the routing table, and *export routing policies* to control which routes a routing protocol advertises from the routing table to its neighbors.

Figure 6: Importing and Exporting Routing Policies




To create a routing policy, you must define the following components:

- *Match conditions*—Criteria that a route must match. If a route matches all of the criteria, one or more actions are applied to the route.
- *Actions*—What to do if a route matches. The actions can specify whether to accept or reject the route, control how a series of policies is evaluated, and manipulate the characteristics associated with a route. You can configure one or more actions.

You typically define match conditions and actions within a *term*. Figure 7 on page 23 shows the routing policy components, including the term.

Figure 7: Routing Policy Components



After defining a routing policy, you then apply it to a routing protocol or to the forwarding table.

This section provides more information about creating routing policies:

- Match Conditions on page 23
- Named Match Conditions on page 24
- Actions on page 25
- Terms on page 25
- Routing Policy Application on page 25

Match Conditions

A *match condition* defines the criteria that a route must match. You can define one or more match conditions. If a route matches all match conditions, one or more actions are applied to the route.

Match conditions fall into two categories: standard and extended. In general, the extended match conditions include criteria that are defined separately from the routing policy (AS path regular expressions, communities, and prefix lists) and are more complex than standard match conditions. The extended match conditions provide many powerful capabilities. For more information about them, see “Extended Match Conditions Configuration” on page 91. The standard match conditions include criteria that are defined within a routing policy and are less complex than the extended match conditions, also called named match conditions.

Table 12 on page 23 describes each match condition, including its category, when you typically use it, and any relevant notes about it. For more information about match conditions, see Table 14 on page 42.

Table 12: Match Conditions

Match Condition	Category	When to Use	Notes
AS path regular expression—A combination of AS numbers and regular expression operators.	Extended	(BGP only) Match a route based on its AS path. (An AS path consists of the AS numbers of all routers a packet must go through to reach a destination.) You can specify an exact match with a particular AS path or a less precise match.	You use regular expressions to match the AS path.

Table 12: Match Conditions *(continued)*

Match Condition	Category	When to Use	Notes
Community—A group of destinations that share a property. (Community information is included as a path attribute in BGP update messages.)	Extended	Match a group of destinations that share a property. Use a routing policy to define a community that specifies a group of destinations you want to match and one or more actions that you want taken on this community.	<p>Actions can be performed on the entire group.</p> <p>You can create multiple communities associated with a particular destination.</p> <p>You can create match conditions using regular expressions.</p>
Prefix list—A named list of IP addresses.	Extended	Match a route based on prefix information. You can specify an exact match of a particular route only.	You can specify a common action only for all prefixes in the list.
Route list—A list of destination prefixes.	Extended	Match a route based on prefix information. You can specify an exact match of a particular route or a less precise match.	You can specify an action for each prefix in the route list or a common action for all prefixes in the route list.
Standard—A collection of criteria that can match a route.	Standard	<p>Match a route based on one of the following criteria: area ID, color, external route, family, instance (routing), interface name, level number, local preference, metric, neighbor address, next-hop address, origin, preference, protocol, routing table name, or tag.</p> <p>For the protocol criterion, you can specify one of the following: BGP, direct, DVMRP, IS-IS, local, MPLS, OSPF, PIM dense mode, PIM sparse mode, RIP, RIPng, static, and aggregate.</p>	None.
Subroutine—A routing policy that is called repeatedly from another routing policy.	Extended	Use an effective routing policy in other routing policies. You can create a subroutine that you can call over and over from other routing policies.	The subroutine action influences but does not necessarily determine the final action. For more information, see “How a Routing Policy Subroutine Is Evaluated” on page 30.

Named Match Conditions

Some match conditions are defined separately from the routing policy and are given names. You then reference the name of the match condition in the definition of the routing policy itself. Named match conditions allow you to do the following:

- Reuse match conditions in other routing policies.
- Read configurations that include complex match conditions more easily.

Named match conditions include communities, prefix lists, and AS path regular expressions. For more information about these match conditions, see Table 12 on page 23.

Actions

An *action* is what the policy framework software does if a route matches all criteria defined in a match condition. You can configure one or more actions in a term. The policy framework software supports the following types of actions:

- Flow control actions, which affect whether to accept or reject the route or whether to evaluate the next term or routing policy
- Actions that manipulate route characteristics
- Trace action, which logs route matches

Manipulating the route characteristics allows you to control which route is selected as the active route to reach a destination. In general, the active route is also advertised to a routing platform's neighbors. You can manipulate the following route characteristics: AS path, class, color, community, damping parameters, destination class, external type, next hop, load balance, local preference, metric, origin, preference, and tag.

For the numeric information (color, local preference, metric, preference, and tag), you can set a specific value or change the value by adding or subtracting a specified amount. The addition and subtraction operations do not allow the value to exceed a maximum value and drop below a minimum value.

For more information about the properties you can change and the addition and subtraction operations, see Table 16 on page 49.

Terms

A *term* is a named structure in which match conditions and actions are defined. You can define one or more terms.

In general, the policy framework software compares a route against the match conditions in the first term in the first routing policy, then goes on to the next term and the next policy if present, and so on, until an explicitly configured or default action of accept or reject is taken. Therefore, the order in which you arrange terms in a policy is relevant.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

Routing Policy Application

After defining a routing policy, as discussed in “Match Conditions” on page 23 and “Actions” on page 25, you can apply it to one of the following:

- Routing protocols—BGP, DVMRP, IS-IS, LDP, MPLS, OSPF, PIM dense mode, PIM sparse mode, PIM sparse-dense mode, RIP, and RIPng

- Pseudoprotocol—Explicitly created routes, which include aggregate and generated routes
- Forwarding table

The following sections discuss the following topics:

- Routing Protocols on page 26
- Routing Policy Application to Routing Protocols on page 27
- Forwarding Table on page 28

Routing Protocols

When applying routing policies to routing protocols, you must know whether each protocol supports import and export policies and the level at which you can apply these policies. Table 13 on page 27 summarizes the import and export policy support for each routing protocol. Table 9 on page 18 also lists explicitly configured routes, which for the purposes of this table are considered a pseudoprotocol. Explicitly configured routes include aggregate and generated routes.

You can apply an import policy to aggregate and generated routes, but you cannot apply an export policy to these routes. These routes cannot be exported from the routing table to the pseudoprotocol, because this protocol is not a real routing protocol. However, aggregate and generated routes can be exported from the routing table to routing protocols.

You cannot apply import policies to the link-state protocols IS-IS and OSPF. As link-state protocols, IS-IS and OSPF exchange routes between systems within an AS. All routers and systems within an AS must share the same link-state database, which includes routes to reachable prefixes and the metrics associated with the prefixes. If an import policy is configured and applied to IS-IS or OSPF, some routes might not be learned or advertised, or the metrics for learned routes might be altered, which would make a consistent link-state database impossible.

For BGP only, you can also apply import and export policies at group and peer levels as well as at the global level. A peer import or export policy overrides a group import or export policy. A group import or export policy overrides a global import or export policy.

For example, if you define an import policy for an individual peer at the peer level and also define an import policy for the group to which it belongs, the import policy defined for the peer level only is invoked. The group import policy is not used for that peer, but it is applied to other peers in that group.

For RIP and RIPng only, you can apply import policies at the global and neighbor levels and export policies at a group level. For more information about RIP and RIPng, see the *JUNOS Routing Protocols Configuration Guide*.

For information about the routing protocols from which the routing table can import routes and to which routing protocols the routing table can export routes, see Table 9 on page 18. For information about the default routing policies for each routing protocol, see Table 11 on page 19.

Table 13: Apply Routing Policies to Protocols

Protocol	Import Policy	Export Policy	Supported Levels
BGP	Yes	Yes	Import: global, group, peer Export: global, group, peer
DVMRP	Yes	Yes	Global
IS-IS	No	Yes	Export: global
LDP	Yes	Yes	Global
MPLS	No	No	–
OSPF	Yes	Yes	Export: global Import: external routes only
PIM dense mode	Yes	Yes	Global
PIM sparse mode	Yes	Yes	Global
Pseudoprotocol—Explicitly configured routes, which include the following: <ul style="list-style-type: none"> ■ Aggregate routes ■ Generated routes 	Yes	No	Import: global
RIP and RIPng	Yes	Yes	Import: global, neighbor Export: group

Routing Policy Application to Routing Protocols

You can apply the following routing policy elements to a routing protocol:

- Routing policy—You can apply a single routing policy to a routing protocol.
- Chain of routing policies—You can apply multiple routing policies (chains) to a routing protocol.
- Policy expression—You can apply a policy expression to a routing protocol. A *policy expression* uses Boolean logical operators with a routing policy and routing policy chains. The logical operators establish rules by which the policy or chains are evaluated.

Forwarding Table

You can apply export policies to routes being exported from the routing table into the forwarding table for the following features:

- Per-packet load balancing
- CoS

For more information about per-packet load balancing, see “Configuring the Load-Balance Per-Packet Action” on page 136. For more information about CoS, see the *JUNOS Class of Service Configuration Guide*.

Evaluating a Routing Policy

This section provides information about how routing policies are evaluated. It discusses the following topics:

- How a Routing Policy Is Evaluated on page 28
- How a Routing Policy Chain Is Evaluated on page 29
- How a Routing Policy Expression Is Evaluated on page 30
- How a Routing Policy Subroutine Is Evaluated on page 30

For specific information about how the various match conditions are evaluated, see “Match Conditions” on page 41 and “Extended Match Conditions Configuration” on page 91.

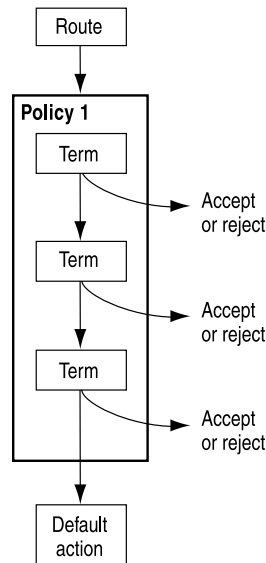
How a Routing Policy Is Evaluated

Figure 8 on page 29 shows how a single routing policy is evaluated. This routing policy consists of multiple terms. Each term consists of match conditions and actions to apply to matching routes. Each route is evaluated against the policy as follows:

1. The route is evaluated against the first term. If it matches, the specified action is taken. If the action is to accept or reject the route, that action is taken and the evaluation of the route ends. If the next term action is specified, if no action is specified, or if the route does not match, the evaluation continues as described in Step 2. If the next policy action is specified, any accept or reject action specified in this term is skipped, all remaining terms in this policy are skipped, all other actions are taken, and the evaluation continues as described in Step 3.
2. The route is evaluated against the second term. If it matches, the specified action is taken. If the action is to accept or reject the route, that action is taken and the evaluation of the route ends. If the next term action is specified, if no action is specified, or if the route does not match, the evaluation continues in a similar manner against the last term. If the next policy action is specified, any accept or reject action specified in this term is skipped, all remaining terms in this policy are skipped, all other actions are taken, and the evaluation continues as described in Step 3.
3. If the route matches no terms in the routing policy or the next policy action is specified, the accept or reject action specified by the default policy is taken. For

more information about the default routing policies, see “Default Routing Policies and Actions” on page 19.

Figure 8: Routing Policy Evaluation

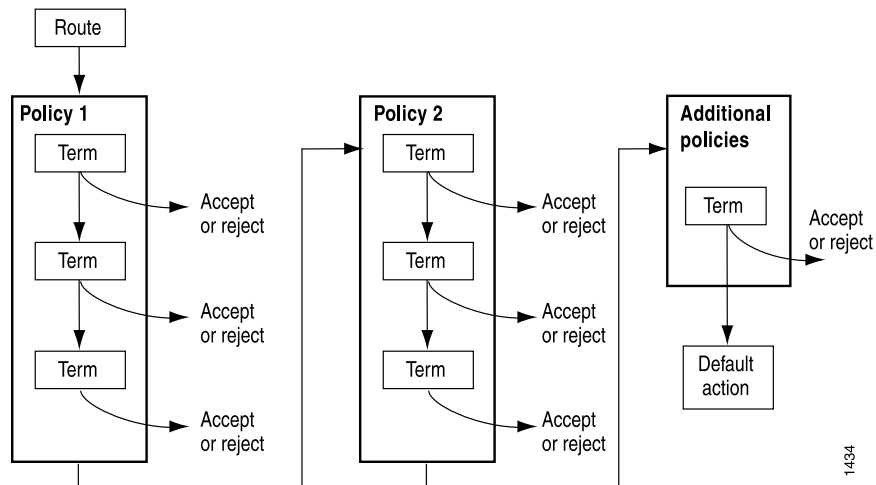


How a Routing Policy Chain Is Evaluated

Figure 9 on page 30 shows how a chain of routing policies is evaluated. These routing policies consist of multiple terms. Each term consists of match conditions and actions to apply to matching routes. Each route is evaluated against the policies as follows:

1. The route is evaluated against the first term in the first routing policy. If it matches, the specified action is taken. If the action is to accept or reject the route, that action is taken and the evaluation of the route ends. If the next term action is specified, if no action is specified, or if the route does not match, the evaluation continues as described in Step 2. If the next policy action is specified, any accept or reject action specified in this term is skipped, all remaining terms in this policy are skipped, all other actions are taken, and the evaluation continues as described in Step 3.
2. The route is evaluated against the second term in the first routing policy. If it matches, the specified action is taken. If the action is to accept or reject the route, that action is taken and the evaluation of the route ends. If the next term action is specified, if no action is specified, or if the route does not match, the evaluation continues in a similar manner against the last term in the first routing policy. If the next policy action is specified, any accept or reject action specified in this term is skipped, all remaining terms in this policy are skipped, all other actions are taken, and the evaluation continues as described in Step 3.

3. If the route does not match a term or matches a term with a next policy action in the first routing policy, it is evaluated against the first term in the second routing policy.
4. The evaluation continues until the route matches a term with an accept or reject action defined or until there are no more routing policies to evaluate. If there are no more routing policies, then the accept or reject action specified by the default policy is taken. For more information about default routing policies, see “Default Routing Policies and Actions” on page 19.

Figure 9: Routing Policy Chain Evaluation

How a Routing Policy Expression Is Evaluated

To understand how a policy expression is evaluated, you must first understand the Boolean logical operators and the associated logic used in evaluating a policy expression. For more information about policy expressions, including how they are evaluated, see “Applying Policy Expressions” on page 59.

How a Routing Policy Subroutine Is Evaluated

Figure 10 on page 32 shows how a subroutine is evaluated. The subroutine is included in the first term of the first routing policy in a chain. Each route is evaluated against the subroutine as follows:

1. The route is evaluated against the first term in the first routing policy. If the route does not match all match conditions specified before the subroutine, the subroutine is skipped and the next term in the routing policy is evaluated (see Step 2). If the route matches all match conditions specified before the subroutine, the route is evaluated against the subroutine. If the route matches the match conditions in any of the subroutine terms, two levels of evaluation occur in the following order:
 - a. The actions in the subroutine term are evaluated. If one of the actions is accept, evaluation of the subroutine ends and a Boolean value of TRUE is returned to the calling policy. If one of the actions is reject, evaluation of the

subroutine ends and FALSE is returned to the calling policy. If one of the actions is meant to manipulate route characteristics, the characteristic is changed regardless of whether accept, reject, or neither action is specified.

If the subroutine does not specify the **accept**, **reject** or **next-policy** action, it uses the **accept** or **reject** action specified by the default policy, and the values of TRUE or FALSE are returned to the calling policy as described in the previous paragraph. (For information about what happens if a termination action is not specified in the term, see “Termination Actions” on page 124. For more information about the default routing policies, see “Default Routing Policies and Actions” on page 19.)

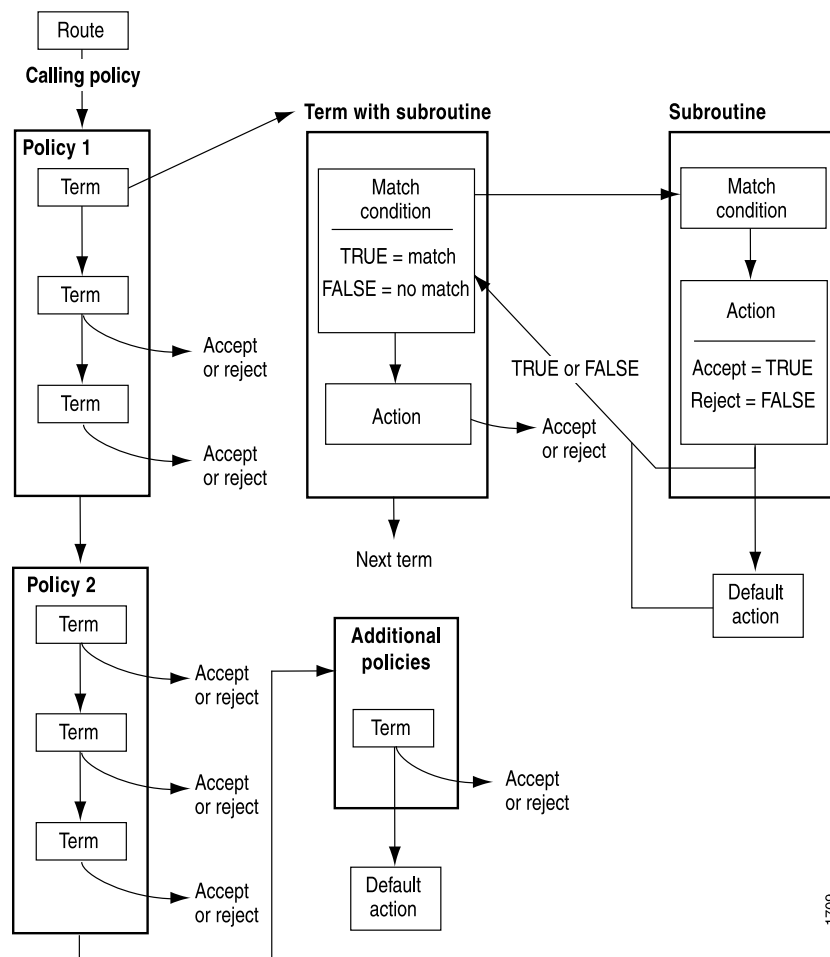
- b. The calling policy’s subroutine match condition is evaluated. During this part of the evaluation, TRUE equals a match and FALSE equals no match. If the subroutine returns TRUE to the calling policy, then the evaluation of the calling policy continues. If the subroutine returns FALSE to the calling policy, then the evaluation of the current term ends and the next term is evaluated.
2. The route is evaluated against the second term in the first routing policy. For information about how the subsequent terms and policies are evaluated, see “How a Routing Policy Chain Is Evaluated” on page 29.



NOTE: If you specify a policy chain as a subroutine, the entire chain acts as a single subroutine. As with other chains, the action specified by the default policy is taken only when the entire chain does not accept or reject a route.



NOTE: If a term defines multiple match conditions, including a subroutine, and a route does not match a condition specified before the subroutine, the evaluation of the term ends and the subroutine is not called and evaluated. In this situation, an action specified in the subroutine that manipulates a route’s characteristics is not implemented.

Figure 10: Routing Policy Subroutine Evaluation

1709

Routing Policy Tests

After you have created a routing policy, you can use the **test policy** command to ensure that the policy produces the results that you expect before applying the policy in a live environment. This command determines whether the routes specified in your routing policy are accepted or rejected. The default action of the **test policy** command is **accept**.



NOTE: The default policy of the **test policy** command accepts all routes from all protocols. Test output can be misleading when you are evaluating protocol-specific conditions.

For example, if you define a policy for BGP that accepts routes of a specified prefix and apply it to BGP as an export policy, the BGP routes that match the prefix are advertised to the BGP peers. However, if you test the same policy using the **test policy** command, the test output might indicate that non-BGP routes have been accepted.

Supported Standards and Drafts

The JUNOS software supports the following standards and drafts:

- RFC 1997, *BGP Communities Attribute*
- RFC 2439, *BGP Route Flap Damping*

Chapter 3

Routing Policy Configuration Statements

To create a routing policy, you can include the `policy-options` statement in the configuration:

```
policy-options {  
  name regular-expression;  
  as-path-group group-name;  
  community name {  
    invert-match;  
    members [ community-ids ];  
  }  
  condition condition-name {  
    if-route-exists address table table-name;  
  }  
  damping name {  
    disable;  
    half-life minutes;  
    max-suppress minutes;  
    reuse number;  
    suppress number;  
  }  
  policy-statement policy-name {  
    term term-name {  
      from {  
        family;  
        match-conditions;  
        policy subroutine-policy-name;  
        prefix-list name;  
        route-filter destination-prefix match-type <actions>;  
        source-address-filter destination-prefix match-type <actions>;  
      }  
      to {  
        match-conditions;  
        policy subroutine-policy-name;  
      }  
      then actions;  
      default-action (accept | reject);  
    }  
  }  
  prefix-list name {  
    ip-addresses;  
  }  
}  
protocols {
```

```

    protocol-name {
        import [ policy-names ];
        export [ policy-names ];
    }
}

```

This section includes the following minimum configurations:

- Minimum Routing Policy Configuration on page 36
- Minimum Routing Policy Chain Configuration on page 36
- Minimum Subroutine Configuration on page 37

Minimum Routing Policy Configuration

To define and apply a routing policy, you must include at least the following statements at the [edit policy-options] and [edit protocols] hierarchy levels. At the [edit protocols] hierarchy level, you can define one or more policy names.

```

[edit]
policy-options {
    policy-statement policy-name {
        term term-name {
            from {
                family family-name;
                match-conditions;
                prefix-list name;
                route-filter destination-prefix match-type <actions>;
                source-address-filter destination-prefix match-type <actions>;
            }
            to {
                match-conditions;
            }
            then actions;
        }
        prefix-list name {
            ip-addresses;
        }
    }
}
protocols {
    protocol-name {
        import [ policy-names ];
        export [ policy-names ];
    }
}
}

```

Minimum Routing Policy Chain Configuration

To define and apply a routing policy chain, you must include at least the following statements at the [edit policy-options] and [edit protocols] hierarchy levels. At the [edit protocols] hierarchy level, you can define a chain of policy names that are evaluated in order.

```

[edit]
policy-options {
  policy-statement policy-name {
    term term-name {
      from {
        family family-name;
        match-conditions;
        prefix-list name;
        route-filter destination-prefix match-type <actions>;
        source-address-filter destination-prefix match-type <actions>;
      }
      to {
        match-conditions;
      }
      then actions;
    }
  }
  policy-statement policy-name {
    term term-name {
      from {
        family family-name;
        match-conditions;
        prefix-list name;
        route-filter destination-prefix match-type <actions>;
        source-address-filter destination-prefix match-type <actions>;
      }
      to {
        match-conditions;
      }
      then actions;
    }
  }
  prefix-list name {
    ip-addresses;
  }
}
protocols {
  protocol-name {
    import [ policy-names ];
    export [ policy-names ];
  }
}

```

Minimum Subroutine Configuration

To configure a routing policy that calls a subroutine from another routing policy, you must include at least the following statements at the [edit *policy-options*] and [edit *protocols*] hierarchy levels. At the [edit *protocols*] hierarchy level, you can define one or more policy names.

```

[edit]
policy-options {
  policy-statement subroutine-policy-name {
    term term-name {

```

```

    from {
        family family-name;
        match-conditions;
        prefix-list name;
        route-filter destination-prefix match-type <actions>;
        source-address-filter destination-prefix match-type <actions>;
    }
    to {
        match-conditions;
    }
    then actions;
}
}
policy-statement policy-name {
    term term-name {
        from {
            family family-name;
            policy subroutine-policy-name;
        }
        to {
            policy subroutine-policy-name;
        }
        then actions;
    }
}
}
protocols {
    protocol-name {
        import [ policy-names ];
        export [ policy-names ];
    }
}

```

Chapter 4

Routing Policy Configuration

This chapter describes the following tasks for configuring routing policies and provides the following examples:

- Defining Routing Policies on page 39
- Applying Routing Policies on page 58
- Examples: Routing Policy Configuration on page 67
- Example: ISP Network Case Study on page 75
- Configuring the Discard Interface on page 87
- Testing Routing Policies on page 88

Defining Routing Policies

To define a routing policy, include the `policy-statement` statement:

```
policy-statement policy-name {  
  term term-name {  
    from {  
      family family-name  
      match-conditions;  
      policy subroutine-policy-name;  
      prefix-list name;  
      route-filter destination-prefix match-type <actions>;  
      source-address-filter destination-prefix match-type <actions>;  
    }  
    to {  
      match-conditions;  
      policy subroutine-policy-name;  
    }  
    then actions;  
  }  
}
```

You can include this statement at the following hierarchy levels:

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

The following sections explain the components of the policy statement and provide configuration examples:

- Routing Policy Name on page 40
- Terms on page 40
- Match Conditions on page 41
- Actions on page 46
- Examples: Defining Routing Policies on page 56

Routing Policy Name

Each routing policy is introduced by the keyword **policy-statement** and a name that identifies it:

```
[edit]
policy-options {
  policy-statement policy-name {
    ...
  }
}
```

The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").

Each routing policy name must be unique within a configuration.

Terms

A routing policy statement consists of one or more named terms:

```
[edit]
policy-options {
  policy-statement policy-name {
    term term-name {
      ...
    }
  }
}
```

The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").

Each term name must be unique within a routing policy.

A policy statement can include one unnamed term. To configure an unnamed term, omit the **term** statement when defining match conditions and actions. The unnamed term is always the last term in the policy and cannot be moved.

Although you can use one unnamed term in each policy statement, we recommend that you name all terms.

Match Conditions

Each term can consist of two statements, **from** and **to**, which define match conditions:

```
[edit]
policy-options {
  policy-statement policy-name {
    term term-name {
      from {
        family family-name;
        match-conditions;
        policy subroutine-policy-name;
        prefix-list name;
        route-filter destination-prefix match-type <actions>;
        source-address-filter source-prefix match-type <actions>;
      }
      to {
        match-conditions;
        policy subroutine-policy-name;
      }
    }
  }
}
```

In the **from** statement, you define the criteria that an incoming route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur.

The **from** statement is optional. If you omit the **from** and the **to** statements, all routes are considered to match.



NOTE: In export policies, omitting the **from** statement from a routing policy term might lead to unexpected results. For more information, see “Side Effects of Omitting the “from” Statement from an Export Policy” on page 64.

In the **to** statement, you define the criteria that an outgoing route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur. You can specify most of the same match conditions in the **to** statement that you can in the **from** statement. In most cases, specifying a match condition in the **to** statement produces the same result as specifying the same match condition in the **from** statement.

The **to** statement is optional. If you omit both the **to** and the **from** statements, all routes are considered to match.



NOTE: All conditions in the **from** and **to** statements must match for the action to be taken. The match conditions are effectively a logical AND operation. Matching in prefix lists and route lists is handled differently. For more information about these match conditions, including how they are evaluated, see “Configuring Prefix Lists” on page 110 and “Configuring Route Lists” on page 114.

Table 14 on page 42 describes the match conditions available for matching an incoming or outgoing route. The table indicates whether you can use the match condition in both **from** and **to** statements and whether the match condition functions the same or differently when used with both statements. If a match condition functions differently in a **from** statement than in a **to** statement, or if the condition cannot be used in one type of statement, there is a separate description for each type of statement. Otherwise, the same description applies to both types of statements.

Table 14 on page 42 also indicates whether the match condition is standard or extended. In general, the extended match conditions include criteria that are defined separately from the routing policy (autonomous system [AS] path regular expressions, communities, and prefix lists) and are more complex than standard match conditions. The extended match conditions provide many powerful capabilities. For more information about them, see “Extended Match Conditions Configuration” on page 91. The standard match conditions include criteria that are defined within a routing policy and are less complex than the extended match conditions.

For examples of how to use the **from** and **to** statements, see “Examples: Routing Policy Configuration” on page 67.

Table 14: Routing Policy Match Conditions

Match Condition	Match Condition Category	from Statement Description	to Statement Description
aggregate-contributor	Standard	Match routes that are contributing to a configured aggregate. This match condition can be used to suppress a contributor in an aggregate route.	
area <i>area-id</i>	Standard	(Open Shortest Path First [OSPF] only) Area identifier. In a from statement used with an export policy, match a route learned from the specified OSPF area when exporting OSPF routes into other protocols.	
as-path <i>name</i>	Extended	(Border Gateway Protocol [BGP] only) Name of an AS path regular expression. For more information, see “Configuring AS Path Regular Expressions” on page 91.	
as-path-group <i>group-name</i>	Extended	(BGP only) Name of an AS path group regular expression. For more information, see “Configuring AS Path Regular Expressions” on page 91.	
color <i>preference</i> color2 <i>preference</i>	Standard	Color value. You can specify preference values (color and color2) that are finer-grained than those specified in the preference and preference2 match conditions. The color value can be a number in the range from 0 through 4,294,967,295 ($2^{32} - 1$). A lower number indicates a more preferred route. For more information about preference values, see the <i>JUNOS Routing Protocols Configuration Guide</i> .	
community [<i>names</i>]	Extended	Name of one or more communities. If you list more than one name, only one name needs to match for a match to occur (the matching is effectively a logical OR operation). For more information, see “Configuring Communities” on page 97.	
external [<i>type</i> <i>metric-type</i>]	Standard	(OSPF only) Match external routes, including routes exported from one level to another. type is an optional keyword. metric can either be 1 or 2. When you do not specify type , this condition matches all external routes. When you specify type , this condition matches only OSPF routes with the specified OSPF metric type.	

Table 14: Routing Policy Match Conditions (*continued*)

Match Condition	Match Condition Category	from Statement Description	to Statement Description
family <i>family-name</i>	Standard	Name of an address family. <i>family-name</i> can be either <i>inet</i> or <i>inet6</i> . Match the address family IP version 4 (IPv4) or IP version 6 (IPv6) of the route. Default setting is <i>inet</i> .	
instance <i>instance-name</i>	Standard	Name of one or more routing instances. Match a route learned from one of the specified instances.	Name of one or more routing instances. Match a route to be advertised over one of the specified instances.
interface <i>interface-name</i>	Standard	Name or IP address of one or more router interfaces. Do not use this qualifier with protocols that are not interface-specific, such as internal BGP (IBGP). Match a route learned from one of the specified interfaces. Direct routes match routes configured on the specified interface.	Name or IP address of one or more router interfaces. Do not use this qualifier with protocols that are not interface-specific, such as IBGP. Match a route to be advertised from one of the specified interfaces.
internal	Standard	Match a routing policy against the internal flag for simplified next-hop self policies.	
level <i>level</i>	Standard	(Intermediate System-to-Intermediate System [IS-IS] only) IS-IS level. Match a route learned from a specified level.	(IS-IS only) IS-IS level. Match a route to be advertised to a specified level.
local-preference <i>value</i>	Standard	(BGP only) BGP local preference (LOCAL_PREF) attribute. The preference value can be a number in the range 0 through 4,294,967,295 ($2^{32} - 1$).	
metric <i>metric</i> <i>metric2</i> <i>metric</i> <i>metric3</i> <i>metric</i> <i>metric4</i> <i>metric</i>	Standard	Metric value. You can specify up to four metric values, starting with <i>metric</i> (for the first metric value) and continuing with <i>metric2</i> , <i>metric3</i> , and <i>metric4</i> . (BGP only) <i>metric</i> corresponds to the multiple exit discriminator (MED), and <i>metric2</i> corresponds to the interior gateway protocol (IGP) metric if the BGP next hop runs back through another route.	

Table 14: Routing Policy Match Conditions (*continued*)

Match Condition	Match Condition Category	from Statement Description	to Statement Description
multicast-scoping (<i>scoping-name</i> <i>number</i>) < (orhigher orlower) >	Standard	<p>Multicast scope value of IPv4 or IPv6 multicast group address. The multicast-scoping name corresponds to an IPv4 prefix. You can match on a specific multicast-scoping prefix or on a range of prefixes. Specify orhigher to match on a scope and numerically higher scopes, or orlower to match on a scope and numerically lower scopes. For more information, see the <i>JUNOS Multicast Protocols Configuration Guide</i>.</p> <p>You can apply this scoping policy to the routing table by including the scope-policy statement at the [edit routing-options] hierarchy level.</p> <p>The <i>number</i> value can be any hexadecimal number from 0 through F. The multicast-scope value is a number from 0 through 15, or one of the following keywords with the associated meanings:</p> <ul style="list-style-type: none"> ■ node-local (value = 1)—No corresponding prefix ■ link-local (value = 2)—Corresponding prefix 224.0.0.0/24 ■ site-local (value = 5)—No corresponding prefix ■ global (value = 14)—Corresponding prefix 224.0.1.0 through 238.255.255.255 ■ organization-local (value = 8)—Corresponding prefix 239.192.0.0/14 	
neighbor <i>address</i>	Standard	<p>Address of one or more neighbors (peers).</p> <p>For BGP, the address can be a directly connected or indirectly connected peer.</p> <p>For all other protocols, the address is the neighbor from which the advertisement is received.</p> <p>NOTE: The neighbor address match condition is not valid for the Routing Information Protocol (RIP).</p>	<p>Address of one or more neighbors (peers).</p> <p>For BGP import policies, specifying to neighbor produces the same result as specifying from neighbor.</p> <p>For BGP export policies, specifying the neighbor match condition has no effect and is ignored.</p> <p>For all other protocols, the to statement matches the neighbor to which the advertisement is sent.</p> <p>NOTE: The neighbor address match condition is not valid for the Routing Information Protocol (RIP).</p>
next-hop <i>address</i>	Standard	Next-hop address or addresses specified in the routing information for a particular route. For BGP routes, matches are performed against each protocol next hop.	
next-hop-type merged	Standard	LDP generates next hop based on RSVP and IP next hops available to use combined with the forwarding-class mapping.	You cannot specify this match condition.
origin <i>value</i>	Standard	<p>(BGP only) BGP origin attribute, which is the origin of the AS path information. The value can be one of the following:</p> <ul style="list-style-type: none"> ■ egp—Path information originated in another AS. ■ igp—Path information originated within the local AS. ■ incomplete—Path information was learned by some other means. 	

Table 14: Routing Policy Match Conditions (*continued*)

Match Condition	Match Condition Category	from Statement Description	to Statement Description
policy [<i>policy-name</i>]	Extended	Name of a policy to evaluate as a subroutine. For information about this extended match condition, see “Configuring Subroutines” on page 123.	
preference <i>preference</i> preference2 <i>preference</i>	Standard	Preference value. You can specify a primary preference value (preference) and a secondary preference value (preference2). The preference value can be a number from 0 through 4,294,967,295 ($2^{32} - 1$). A lower number indicates a more preferred route. To specify even finer-grained preference values, see the color and color2 match conditions in this table. For more information about preference values, see the <i>JUNOS Routing Protocols Configuration Guide</i> .	
prefix-list <i>prefix-list-name</i> <i>ip-addresses</i>	Extended	Named list of IP addresses. You can specify an exact match with incoming routes. For information about this extended match condition, see “Configuring Prefix Lists” on page 110.	You cannot specify this match condition.
prefix-list-filter <i>prefix-list-name</i> <i>match-type</i>	Extended	Named prefix list. You can specify prefix length qualifiers for the list of prefixes in the prefix list. For information about this extended match condition, see “Configuring Prefix Lists” on page 110.	You cannot specify this match condition.
protocol <i>protocol</i>	Standard	Name of the protocol from which the route was learned or to which the route is being advertised. It can be one of the following: access , access-internal , aggregate , bgp , direct , dvmrp , isis , local , ospf , ospf2 , ospf3pim-dense , pim-sparse , rip , ripng , or static . NOTE: The ospf2 statement matches on OSPFv2 routes. The ospf3 statement matches on OSPFv3 routes. The ospf statement matches on both OSPFv2 and OSPFv3 routes. For more information about access routes and access-internal routes, see “Importing and Exporting Access and Access-Internal Routes in a Routing Policy” on page 57.	
rib <i>routing-table</i>	Standard	Name of a routing table. The value of <i>routing-table</i> can be one of the following: <ul style="list-style-type: none"> ■ inet.0—Unicast IPv4 routes ■ <i>instance-name</i> inet.0—Unicast IPv4 routes for a particular routing instance ■ inet.1—Multicast IPv4 routes ■ inet.2—Unicast IPv4 routes for multicast reverse-path forwarding (RPF) lookup ■ inet.3—Multiprotocol Label Switching (MPLS) routes ■ mpls.0—MPLS routes for label-switched path (LSP) next hops ■ inet6.0—Unicast IPv6 routes 	

Table 14: Routing Policy Match Conditions (*continued*)

Match Condition	Match Condition Category	from Statement Description	to Statement Description
<code>route-filter</code> <code>destination-prefix</code> <code>match-type <actions></code>	Extended	List of destination prefixes. When specifying a destination prefix, you can specify an exact match with a specific route or a less precise match using match types. You can configure either a common action that applies to the entire list or an action associated with each prefix. For more information, see “Configuring Route Lists” on page 114.	You cannot specify this match condition.
<code>route-type value</code>	Standard	Type of route. The value can be one of the following: <ul style="list-style-type: none"> ■ <code>external</code>—External route. ■ <code>internal</code>—Internal route. 	
<code>source-address-filter</code> <code>destination-prefix</code> <code>match-type <actions></code>	Extended	List of multicast source addresses. When specifying a source address, you can specify an exact match with a specific route or a less precise match using match types. You can configure either a common action that applies to the entire list or an action associated with each prefix. For more information, see “Configuring Route Lists” on page 114.	You cannot specify this match condition.
<code>state (active inactive)</code>	Standard	(BGP export only) Match on the following types of advertised routes: <ul style="list-style-type: none"> ■ <code>active</code>—An active BGP route ■ <code>inactive</code>—A route advertised to internal BGP peers as the best external path even if the best path is an internal route ■ <code>inactive</code>—A route advertised by BGP as the best route even if the routing table did not select it to be an active route 	
<code>tag string tag2 string</code>	Standard	Tag value. You can specify two tag strings: <code>tag</code> (for the first string) and <code>tag2</code> . These values are local to the router and can be set on configured routes or by using an import routing policy. You can specify multiple tags under one match condition by including the tags within a bracketed list. For example: <code>from tag [tag1 tag2 tag3]</code> ; For OSPF and IS-IS, the <code>tag</code> match conditions match the 32-bit tag field in external link-state advertisement (LSA) packets.	

Actions

Each term can include a `then` statement, which defines the actions to take if a route matches all the conditions in the `from` and `to` statements:

```
[edit]
policy-options {
  policy-statement policy-name {
```

```

term term-name {
  from {
    family family-name;
    match-conditions;
    policy subroutine-policy-name;
    prefix-list name;
    route-filter destination-prefix match-type <actions>;
    source-address-filter destination-prefix match-type <actions>;
  }
  to {
    match-conditions;
    policy subroutine-policy-name;
  }
  then {
    actions;
  }
}
}

```

If a term does not have **from** and **to** statements, all routes are considered to match, and the actions apply to them all.

The **then** statement is optional.

You can specify one or more actions. There are three types of actions:

- Flow control actions, which affect whether to accept or reject the route and whether to evaluate the next term or routing policy.
- Actions that manipulate route characteristics.
- Trace action, which logs route matches.



NOTE: When you specify an action that manipulates the route characteristics, the changes occur in a copy of the source route. The source route itself does not change. The effect of the action is visible only after the route is imported into or exported from the routing table. To view the source route before the routing policy has been applied, use the **show route receive-protocol** command. To view a route after an export policy has been applied, use the **show route advertised-protocol** command.

During policy evaluation, the characteristics in the copy of the source route always change immediately after the action is evaluated. However, the route is not copied to the routing table or a routing protocol until the completion of the policy evaluation is complete.

If you do not include a **then** statement, one of the following occurs:

- The next term in the routing policy, if one is present, is evaluated.
- If there are no more terms in the routing policy, the next routing policy, if one is present, is evaluated.

- If there are no more terms or routing policies, the accept or reject action specified by the default policy is taken. For more information, see “Default Routing Policies and Actions” on page 19.

The following sections discuss the following action attributes:

- Flow Control Actions on page 48
- Actions That Manipulate Route Characteristics on page 49
- Trace Action on page 53
- Final Action on page 54
- Default Action on page 54
- Example: Configure the Default Action for a Policy on page 55
- Route List Actions on page 56

Flow Control Actions

Table 15 on page 48 lists the flow control actions. You can specify one of these actions along with the trace action (see “Trace Action” on page 53) or one or more of the actions that manipulate route characteristics (see “Actions That Manipulate Route Characteristics” on page 49).

Table 15: Flow Control Actions

Flow Control Action	Description
accept	Accept the route and propagate it. After a route is accepted, no other terms in the routing policy and no other routing policies are evaluated.
default-action accept	Accept and override any action intrinsic to the protocol. This is a nonterminating policy action.
reject	Reject the route and do not propagate it. After a route is rejected, no other terms in the routing policy and no other routing policies are evaluated.
default-action reject	Reject and override any action intrinsic to the protocol. This is a nonterminating policy action.
next term	<p>Skip to and evaluate the next term in the same routing policy. Any accept or reject action specified in the then statement is skipped. Any actions in the then statement that manipulate route characteristics are applied to the route.</p> <p>next term is the default control action if a match occurs and you do not specify a flow control action.</p>
next policy	<p>Skip to and evaluate the next routing policy. Any accept or reject action specified in the then statement is skipped. Any actions in the then statement that manipulate route characteristics are applied to the route.</p> <p>next policy is the default control action if a match occurs, you do not specify a flow control action, and there are no further terms in the current routing policy.</p>

Actions That Manipulate Route Characteristics

You can specify one or more of the actions listed in Table 16 on page 49 to manipulate route characteristics.

Table 16: Actions That Manipulate Route Characteristics

Action	Description
<code>as-path-prepend as-path</code>	<p>(BGP only) Affix one or more AS numbers at the beginning of the AS path. If specifying more than one AS number, include the numbers in quotation marks. The AS numbers are added after the local AS number has been added to the path. This action adds AS numbers to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the affixed AS numbers are placed within a confederation sequence. Otherwise, the affixed AS numbers are placed with a nonconfederation sequence. For more information, see “Configuring the AS Path Prepend Action” on page 129.</p> <p>Beginning with JUNOS Release 9.1, the range that you can configure for AS numbers has been extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for 4-octet AS Number Space</i>. The JUNOS software continues to support 2-byte AS numbers. For more information about configuring the AS number for the router and BGP, see the <i>JUNOS Routing Protocols Configuration Guide</i>.</p>
<code>as-path-expand last-as count n</code>	<p>(BGP only) Extract the last AS number in the existing AS path and affix that AS number to the beginning of the AS path <i>n</i> times, where <i>n</i> is a number from 1 through 32. The AS number is added before the local AS number has been added to the path. This action adds AS numbers to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the affixed AS numbers are placed within a confederation sequence. Otherwise, the affixed AS numbers are placed within a nonconfederation sequence. This option is typically used in non-IBGP export policies.</p>
<code>class class-name</code>	<p>(Class of service [CoS] only) Apply the specified class-of-service parameters to routes installed into the routing table. For more information, see the <i>JUNOS Class of Service Configuration Guide</i>.</p>
<code>color preference color2 preference</code>	<p>Set the preference value to the specified value. The <code>color</code> and <code>color2</code> preference values are even more fine-grained than those specified in the <code>preference</code> and <code>preference2</code> actions. The <code>color</code> value can be a number in the range from 0 through 4,294,967,295 ($2^{32} - 1$). A lower number indicates a more preferred route.</p> <p>If you set the preference with the <code>color</code> action, the value is internal to the JUNOS software and is not transitive.</p> <p>For more information about preference values, see the <i>JUNOS Routing Protocols Configuration Guide</i>.</p>
<code>color (add subtract) number</code> <code>color2 (add subtract) number</code>	<p>Change the color preference value by the specified amount. If an addition operation results in a value that is greater than 4,294,967,295 ($2^{32} - 1$), the value is set to $2^{32} - 1$. If a subtraction operation results in a value less than 0, the value is set to 0. If an attribute value is not already set at the time of the addition or subtraction operation, the attribute value defaults to a value of 0 regardless of the amount specified. If you perform an addition to an attribute with a value of 0, the number you add becomes the resulting attribute value.</p>
<code>community (+ add) [names]</code>	<p>(BGP only) Add the specified communities to the set of communities in the route. For more information, see “Configuring Communities” on page 97.</p>
<code>community (– delete) [names]</code>	<p>(BGP only) Delete the specified communities from the set of communities in the route. For more information, see “Configuring Communities” on page 97.</p>

Table 16: Actions That Manipulate Route Characteristics (*continued*)

Action	Description
community (= set) [<i>names</i>]	(BGP only) Replace any communities that were in the route in with the specified communities. For more information, see “Configuring Communities” on page 97.
cos-next-hop-map <i>map-name</i>	Set CoS-based next-hop map in forwarding table.
damping <i>name</i>	<p>(BGP only) Apply the specified route-damping parameters to the route. These parameters override the default damping parameters. This action is useful only in an import policy, because the damping parameters affect the state of routes in the routing table.</p> <p>To apply damping parameters, you must enable BGP flap damping as described in the <i>JUNOS Routing Protocols Configuration Guide</i>, and you must create a named list of parameters as described in “Configuring the Damping Action” on page 130.</p>
destination-class <i>destination-class-name</i>	<p>Maintain packet counts for a route passing through your network, based on the destination address in the packet. You can do the following:</p> <ul style="list-style-type: none"> ■ Configure group destination prefixes by configuring a routing policy; see “Defining Routing Policies” on page 39 and “Examples: Routing Policy Configuration” on page 67. ■ Apply that routing policy to the forwarding table with the corresponding destination class; see “Applying Routing Policies to the Forwarding Table” on page 66. For more information about the forwarding-table configuration statement, see the <i>JUNOS Routing Protocols Configuration Guide</i>. ■ Enable packet counting on one or more interfaces by including the destination-class-usage statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting] hierarchy level (see the <i>JUNOS Class of Service Configuration Guide</i>). See “Examples: Routing Policy Configuration” on page 67. ■ View the output by using one of the following commands: show interfaces destination-class (all <i>destination-class-name logical-interface-name</i>), show interfaces interface-name extensive, or show interfaces interface-name statistics (see the <i>JUNOS Interfaces Command Reference</i>). ■ To configure a packet count based on the source address, use the source-class statement described in this table.
metric expression metric multiplier <i>x</i> offset <i>a</i> metric2 multiplier <i>y</i> offset <i>b</i>	<p>Calculate a metric based on the current values of metric and metric2.</p> <p>This policy action overrides the current value of the metric attribute with the result of the expression</p> $((x * \text{metric}) + a) + ((y * \text{metric2}) + b)$ <p>where metric and metric2 are the current input values. Metric multipliers are limited in range to eight significant digits.</p>
external type <i>metric</i>	Set the external metric type for routes exported by OSPF. You must specify the keyword type .

Table 16: Actions That Manipulate Route Characteristics (*continued*)

Action	Description
forwarding-class forwarding-class-name	<p>Create the forwarding class that includes packets based on both the destination address and the source address in the packet. You can do the following:</p> <ul style="list-style-type: none"> ■ Configure group prefixes by configuring a routing policy; see “Defining Routing Policies” on page 39 and “Examples: Routing Policy Configuration” on page 67. ■ Apply that routing policy to the forwarding table with the corresponding forwarding class; see “Applying Routing Policies to the Forwarding Table” on page 66. For more information about the <code>forwarding-table</code> configuration statement, see the <i>JUNOS Routing Protocols Configuration Guide</i>. ■ Enable packet counting on one or more interfaces by using the procedure described in either the <code>destination-class</code> or <code>source-class</code> actions defined in this table.
install-nexthop <strict> lsp lsp-name	Choose which next hops, among a set of equal LSP next hops, are installed in the forwarding table. Use the export policy for the forwarding table to specify the LSP next hop to be used for the desired routes. Specify the <code>strict</code> option to enable strict mode, which checks to see if any of the LSP next hops specified in the policy are up. If none of the specified LSP next hops are up, the policy installs the discard next hop.
load-balance per-packet	(For export to the forwarding table only) Install all next-hop addresses in the forwarding table and have the forwarding table perform per-packet load balancing. This policy action allows you to optimize VPLS traffic flows across multiple paths. For more information, see “Configuring the Load-Balance Per-Packet Action” on page 136.
local-preference value	(BGP only) Set the BGP local preference (<code>LOCAL_PREF</code>) attribute. The preference value can be a number in the range from 0 through 4,294,967,295 ($2^{32} - 1$).
local-preference (add subtract) number	<p>Change the local preference value by the specified amount. If an addition operation results in a value that is greater than 4,294,967,295 ($2^{32} - 1$), the value is set to $2^{32} - 1$. If a subtraction operation results in a value less than 0, the value is set to 0. If an attribute value is not already set at the time of the addition or subtraction operation, the attribute value defaults to a value of 0 regardless of the amount specified. If you perform an addition to an attribute with a value of 0, the number you add becomes the resulting attribute value.</p> <p>For BGP, if the attribute value is not known, it is initialized to 100 before the routing policy is applied.</p>
metric metric metric2 metric metric3 metric metric4 metric	<p>Set the metric. You can specify up to four metric values, starting with <code>metric</code> (for the first metric value) and continuing with <code>metric2</code>, <code>metric3</code>, and <code>metric4</code>.</p> <p>(BGP only) <code>metric</code> corresponds to the MED, and <code>metric2</code> corresponds to the IGP metric if the BGP next hop loops through another router.</p>
metric (add subtract) number metric2 (add subtract) number metric3 (add subtract) number metric4 (add subtract) number	Change the metric value by the specified amount. If an addition operation results in a value that is greater than 4,294,967,295 ($2^{32} - 1$), the value is set to $2^{32} - 1$. If a subtraction operation results in a value less than 0, the value is set to 0. If an attribute value is not already set at the time of the addition or subtraction operation, the attribute value defaults to a value of 0 regardless of the amount specified. If you perform an addition to an attribute with a value of 0, the number you add becomes the resulting attribute value.
metric (igp minimum-igp) site-offset	(BGP only) Change the metric (MED) value by the specified negative or positive offset. This action is useful only in an external BGP (EBGP) export policy.

Table 16: Actions That Manipulate Route Characteristics (*continued*)

Action	Description
next-hop <i>address</i> (discard next-table peer-address reject self)	<p>Set the next hop. When the advertising protocol is BGP, you can set the next hop only when any third-party next hop can be advertised; that is, when you are using IBGP or EBGP confederations.</p> <p>If you specify self, the next-hop address is replaced by one of the local router's addresses. The advertising protocol determines which address to use. When the advertising protocol is BGP, this address is set to the local IP address used for the BGP adjacency. A router cannot install routes with itself as the next hop.</p> <p>If you specify peer-address, the next-hop address is replaced by the peer's IP address. This option is valid only in import policies. Primarily used by BGP to enforce using the peer's IP address for advertised routes, this option is meaningful only when the next hop is the advertising router or another directly connected router.</p> <p>If you specify discard, the next-hop address is replaced by a discard next hop.</p> <p>If you specify next-table, the router performs a forwarding lookup in the specified table.</p> <p>If you specify reject, the next-hop address is replaced by a reject next hop.</p>
origin <i>value</i>	<p>(BGP only) Set the BGP origin attribute to one of the following values:</p> <ul style="list-style-type: none"> ■ igp—Path information originated within the local AS. ■ egp—Path information originated in another AS. ■ incomplete—Path information learned by some other means.
preference <i>preference</i> preference2 <i>preference</i>	<p>Set the preference value. You can specify a primary preference value (preference) and a secondary preference value (preference2). The preference value can be a number in the range from 0 through 4,294,967,295 ($2^{32} - 1$). A lower number indicates a more preferred route.</p> <p>To specify even finer-grained preference values, see the color and color2 actions in this table.</p> <p>If you set the preference with the preference action, the new preference remains associated with the route. The new preference is internal to the JUNOS software and is not transitive.</p> <p>For more information about preference values, see the <i>JUNOS Routing Protocols Configuration Guide</i>.</p>
preference (add subtract) <i>number</i> preference2 (add subtract) <i>number</i>	<p>Change the preference value by the specified amount. If an addition operation results in a value that is greater than 4,294,967,295 ($2^{32} - 1$), the value is set to $2^{32} - 1$. If a subtraction operation results in a value less than 0, the value is set to 0. If an attribute value is not already set at the time of the addition or subtraction operation, the attribute value defaults to a value of 0 regardless of the amount specified. If you perform an addition to an attribute with a value of 0, the number you add becomes the resulting attribute value.</p>
priority (low medium high)	<p>(OSPF import only) Specify a priority for prefixes included in an OSPF import policy. Prefixes learned through OSPF are installed in the routing table based on the priority assigned to the prefixes. Prefixes assigned a priority of high are installed first, while prefixes assigned a priority of low are installed last. For more detailed information about configuring priority for prefixes included in OSPF import policy, see the <i>JUNOS Routing Protocols Configuration Guide</i>.</p> <p>NOTE: OSPF import policy can only be used to set priority or to filter OSPF external routes. If an OSPF import policy is applied that results in a reject terminating action for a nonexternal route, then the reject action is ignored and the route is accepted anyway.</p>

Table 16: Actions That Manipulate Route Characteristics (*continued*)

Action	Description
<code>source-class source-class-name</code>	<p>Maintain packet counts for a route passing through your network, based on the source address. You can do the following:</p> <ul style="list-style-type: none"> ■ Configure group source prefixes by configuring a routing policy; see “Defining Routing Policies” on page 39 and “Examples: Routing Policy Configuration” on page 67. ■ Apply that routing policy to the forwarding table with the corresponding source class; see “Applying Routing Policies to the Forwarding Table” on page 66. For more information about the <code>forwarding-table</code> configuration statement, see the <i>JUNOS Routing Protocols Configuration Guide</i>. ■ Enable packet counting on one or more interfaces by including the <code>source-class-usage interface-name</code> statement at the [edit interfaces <i>logical-unit-number</i> unit family inet accounting] hierarchy level (see the <i>JUNOS Network Interfaces Configuration Guide</i>). Also, follow the <code>source-class-usage</code> statement with the <code>input</code> or <code>output</code> statement to define the inbound and outbound interfaces on which traffic monitored for source-class usage (SCU) is arriving and departing (or define one interface for both). The complete syntax is [edit interfaces <i>interface-name</i> unit family inet accounting source-class-usage (input output input output) <i>unit-number</i>]. See the example in “Examples: Routing Policy Configuration” on page 67. ■ View the output by using one of the following commands: <code>show interfaces interface-name source-class source-class-name</code>, <code>show interfaces interface-name extensive</code>, or <code>show interfaces interface-name statistics</code> (see the <i>JUNOS Interfaces Command Reference</i>). ■ To configure a packet count based on the destination address, use the <code>destination-class</code> statement described in this table. ■ For a detailed source-class usage example configuration, see the <i>JUNOS Feature Guide</i>.
<code>tag tag tag2 tag</code>	<p>Set the tag value. You can specify two tag strings: <code>tag</code> (for the first string) and <code>tag2</code>. These values are local to the router.</p> <p>(For OSPF and IS-IS only) The <code>tag</code> actions set the 32-bit tag field in OSPF external link-state advertisement (LSA) packets and the 32-bit flag in the IS-IS IP prefix type length values (TLV).</p>
<code>tag (add subtract) number tag2 (add subtract) number</code>	<p>Change the tag value by the specified amount. If an addition operation results in a value that is greater than 4,294,967,295 ($2^{32} - 1$), the value is set to $2^{32} - 1$. If a subtraction operation results in a value less than 0, the value is set to 0. If an attribute value is not already set at the time of the addition or subtraction operation, the attribute value defaults to a value of 0 regardless of the amount specified. If you perform an addition to an attribute with a value of 0, the number you add becomes the resulting attribute value.</p>

Trace Action

If you specify the trace action, the match is logged to a trace file. To set up a trace file, you must specify the following elements in the global `traceoptions` statement:

- Trace filename
- policy option in the `flag` statement

For more information about the global `traceoptions` statement, see the *JUNOS Routing Protocols Configuration Guide*.

The following example uses the trace filename of `policy-log`:

```
[edit]
routing-options {
  traceoptions {
    file "policy-log";
    flag policy;
  }
}
```

This action does not affect the flow control during routing policy evaluation.

If a term that specifies a trace action also specifies a flow control action, the name of the term will be logged in the trace file. If a term specifies a trace action only, the word `<default>` only will be logged.

Final Action

In addition to specifying an action using the `then` statement in a named term, you can also specify an action using the `then` statement in an unnamed term, as follows:

```
[edit]
policy-options {
  policy-statement policy-name {
    term term-name {
      from {
        family family-name;
        match-conditions;
        policy subroutine-policy-name;
        prefix-list name;
        route-filter destination-prefix match-type <actions>;
        source-address-filter source-prefix match-type <actions>;
      }
      to {
        match-conditions;
        policy subroutine-policy-name;
      }
      then {
        actions;
      }
    }
  }
  then action;
}
```

Default Action

The `default-action` statement overrides any action intrinsic to the protocol. This action is also nonterminating, so that various policy terms can be evaluated before the policy is terminated. You can specify a default action, either `accept` or `reject`, as follows:

```
[edit]
policy-options {
  policy-statement policy-name {
    term term-name {
```

```

    from {
        family family-name;
        match-conditions;
        policy subroutine-policy-name;
        prefix-list name;
        route-filter destination-prefix match-type <actions>;
        source-address-filter source-prefix match-type <actions>;
    }
    to {
        match-conditions;
        policy subroutine-policy-name;
    }
    then {
        actions;
        default-action (accept | reject);
    }
}
}
}

```

The resulting action is set either by the protocol or by the last policy term that is matched.

Example: Configure the Default Action for a Policy

Configure a routing policy that matches routes based on three policy terms. If the route matches the first term, a certain community tag is attached. If the route matches two separate terms, then both community tags are attached. If the route does not match any terms, it is rejected (protocol's default action). Note that the terms **hub** and **spoke** are mutually exclusive.

```

[edit]
policy-options {
  policy-statement test {
    term set-default {
      then default-action reject;
    }
    term hub {
      from interface ge-2/1/0.5;
      then {
        community add test-01-hub;
        default-action accept;
      }
    }
    term spoke {
      from interface [ ge-2/1/0.1 ge-2/1/0.2 ];
      then {
        community add test-01-spoke;
        default-action accept;
      }
    }
    term management {
      from protocol direct;
      then {
        community add management;
      }
    }
  }
}

```

```

        default-action accept;
    }
}
}

```

Route List Actions

If you specify route lists in the `from` statement, for each route in the list, you can specify an action to take on that individual route directly, without including a `then` statement. For more information, see “Configuring Route Lists” on page 114.

Examples: Defining Routing Policies

This section provides examples of defining routing policies. For more examples, see “Examples: Routing Policy Configuration” on page 67.

Defining a Routing Policy from BGP to IS-IS

Accept BGP routes advertised by the peer `192.168.1.1`. If a route matches, it is accepted, and no further evaluation is performed on that route. If a route does not match, the accept or reject action specified by the default policy is taken. (For more information about the default routing policies, see “Default Routing Policies and Actions” on page 19.) If you apply this routing policy to imported BGP routes, only the routes learned from the peer `192.168.1.1` and BGP transit routes are accepted from BGP peers.

```

[edit]
policy-options {
  policy-statement bgp-to-isis {
    term term1 {
      from {
        neighbor 192.168.1.1;
      }
      then {
        accept;
      }
    }
  }
}

```

Using Routing Policy to Set a Preference

Define a routing policy which matches routes from specific next hops that are being advertised to specific neighbors and which sets a preference. If a route does not match the first term, it is evaluated by the second term. If it still does not match, the next routing policy, if configured, is evaluated; then the accept or reject action specified by the default policy is taken. (For more information about the default routing policies, see “Default Routing Policies and Actions” on page 19.)

```

[edit]
policy-options {
  policy-statement set-preference {

```

```

term term1 {
  from {
    next-hop [ 10.0.0.1 10.0.0.2 ];
  }
  to {
    neighbor 192.168.1.1;
  }
  then {
    preference 10;
  }
}
term term2 {
  from {
    next-hop 10.0.0.3;
  }
  to {
    neighbor 192.168.1.1;
  }
  then {
    preference 15;
  }
}
}

```

Importing and Exporting Access and Access-Internal Routes in a Routing Policy

Configure import and export of access routes and access-internal routes in a routing policy. These routes are used by the DHCP application on a video services router to represent either the end users or the networks behind the attached video services router. (For more information about configuring DHCP relay on the router, see “Configuring the Extended DHCP Relay Agent” on page 299.)

An access route represents a network behind an attached video services router, and is set to a preference of 13. An access-internal route is a /32 route that represents a directly attached end user, and is set to a preference of 12.

```

[edit]
policy-options {
  policy-statement foo {
    term term1 {
      from protocol {
        access;
        access-internal;
      }
      then accept;
    }
  }
}

```

Applying Routing Policies

For a routing policy to take effect, you must apply it to either a routing protocol or the forwarding table. This section contains the following information:

- Applying Routing Policies to a Routing Protocol on page 58
- Applying Routing Policies to the Forwarding Table on page 66
- Examples: Applying Routing Policies on page 66

Applying Routing Policies to a Routing Protocol

Before applying routing policies to routing protocols, you must know if each protocol supports import and export policies and the level at which you can apply these policies. Table 9 on page 18 summarizes the import and export policy support for each routing protocol and the level at which you can apply these policies.

For more information about applying routing policies to individual routing protocols, see the *JUNOS Routing Protocols Configuration Guide*.

The following sections discuss applying routing policies to a routing protocol and some side effects of these tasks:

- Applying a Routing Policy on page 58
- Applying a Routing Policy Chain on page 59
- Applying Policy Expressions on page 59
- Side Effects of Omitting the "from" Statement from an Export Policy on page 64

Applying a Routing Policy

To apply a routing policy to a routing protocol, include the **import** and **export** statements:

```
import [ policy-names ];
export [ policy-names ];
```

You can include the statements at the following hierarchy levels:

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

In the **import** statement, list the name of the routing policy to be evaluated when routes are imported into the routing table from the routing protocol.

In the **export** statement, list the name of the routing policy to be evaluated when routes are being exported from the routing table into a dynamic routing protocol. Only active routes are exported from the routing table.

You can reference the same routing policy one or more times in the same or different **import** and **export** statements.

For information about how the policy framework software evaluates a routing policy, see “How a Routing Policy Is Evaluated” on page 28.

Applying a Routing Policy Chain

To apply multiple routing policies (chains) to a routing protocol, include the **import** and **export** statements:

```
import [ policy-names ];
export [ policy-names ];
```

You can include the statements at the following hierarchy levels:

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

In the **import** statement, list the names of multiple routing policies to be evaluated when routes are imported into the routing table from the routing protocol.

In the **export** statement, list the names of multiple routing policies to be evaluated when routes are being exported from the routing table into a dynamic routing protocol. Only active routes are exported from the routing table.

You can reference the same routing policy one or more times in the same or different **import** and **export** statements.

The policy framework software evaluates the routing policies sequentially, from left to right. If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during the evaluation of the remaining policies. For example, if the action specified in the first policy of a chain sets a route’s metric to 500, this route matches the criterion of **metric 500** defined in the next policy.

For more information about routing policy chain evaluation, see “How a Routing Policy Chain Is Evaluated” on page 29.

Applying Policy Expressions

Policy expressions give the policy framework software a different way to evaluate routing policies. A *policy expression* uses Boolean logical operators with policies. The logical operators establish rules by which the policies are evaluated.

During evaluation of a routing policy in a policy expression, the policy action of accept, reject, or next policy is converted to the value of TRUE or FALSE. This value is then evaluated against the specified logical operator to produce output of either TRUE or FALSE. The output is then converted back to a flow control action of accept, reject, or next policy. The result of the policy expression is applied as it would be applied to a single policy; the route is accepted or rejected and the evaluation ends, or the next policy is evaluated.

Table 17 on page 60 summarizes the policy actions and their corresponding TRUE and FALSE values and flow control action values. Table 18 on page 60 describes the

logical operators. For complete information about policy expression evaluation, see “How a Policy Expression Is Evaluated” on page 62.

You must enclose a policy expression in parentheses. You can place a policy expression anywhere in the `import` or `export` statements and in the `from policy` statement.

Table 17: Policy Action Conversion Values

Policy Action	Conversion Value	Flow Control Action Conversion Value
Accept	TRUE	Accept
Reject	FALSE	Reject
Next policy	TRUE	Next policy

Table 18: Policy Expression Logical Operators

Logical Operator	Policy Expression Logic	How Logical Operator Affects Policy Expression Evaluation
&& (Logical AND)	<p>Logical AND requires that all values must be TRUE to produce output of TRUE.</p> <p>Routing policy value of TRUE and TRUE produces output of TRUE. Value of TRUE and FALSE produces output of FALSE. Value of FALSE and FALSE produces output of FALSE.</p>	<p>If the first routing policy returns the value of TRUE, the next policy is evaluated. If the first policy returns the value of FALSE, the evaluation of the expression ends and subsequent policies in the expression are not evaluated.</p>
(Logical OR)	<p>Logical OR requires that at least one value must be TRUE to produce output of TRUE.</p> <p>Routing policy value of TRUE and FALSE produces output of TRUE. Value of TRUE and TRUE produces output of TRUE. Value of FALSE and FALSE produces output of FALSE.</p>	<p>If the first routing policy returns the value of TRUE, the evaluation of the expression ends and subsequent policies in the expression are not evaluated. If the first policy returns the value of FALSE, the next policy is evaluated.</p>
! (Logical NOT)	<p>Logical NOT reverses value of TRUE to FALSE and of FALSE to TRUE. It also reverses the actions of accept and next policy to reject, and reject to accept.</p>	<p>If used with the logical AND operator and the first routing policy value of FALSE is reversed to TRUE, the next policy is evaluated. If the value of TRUE is reversed to FALSE, the evaluation of the expression ends and subsequent policies in the expression are not evaluated.</p> <p>If used with the logical OR operator and the first routing policy value of FALSE is reversed to TRUE, the evaluation of the expression ends and subsequent policies in the expression are not evaluated. If the value of TRUE is reversed to FALSE, the next policy is evaluated.</p> <p>If used with a policy and the flow control action is accept or next policy, these actions are reversed to reject. If the flow control action is reject, this action is reversed to accept.</p>

For more information, see the following sections:

- Policy Expression Examples on page 61
- How a Policy Expression Is Evaluated on page 62
- Example: Evaluating Policy Expressions on page 63

Policy Expression Examples

The following examples show how to use the logical operators to create policy expressions:

- Logical AND—In the following example, **policy1** is evaluated first. If after **policy1** is evaluated, a value of TRUE is returned, **policy2** is evaluated. If a value of FALSE is returned, **policy2** is not evaluated.

```
export (policy1 && policy2)
```

- Logical OR—In the following example, **policy1** is evaluated first. If after **policy1** is evaluated, a value of TRUE is returned, **policy2** is not evaluated. If a value of FALSE is returned, **policy2** is evaluated.

```
export (policy1 || policy2)
```

- Logical OR and logical AND—In the following example, **policy1** is evaluated first. If after **policy1** is evaluated, a value of TRUE is returned, **policy2** is skipped and **policy3** is evaluated. If after **policy1** is evaluated, a value of FALSE is returned, **policy2** is evaluated. If **policy2** returns a value of TRUE, **policy3** is evaluated. If **policy2** returns a value of FALSE, **policy3** is not evaluated.

```
export [(policy1 || policy2) && policy3]
```

- Logical NOT—In the following example, **policy1** is evaluated first. If after **policy1** is evaluated, a value of TRUE is returned, the value is reversed to FALSE and **policy2** is not evaluated. If a value of FALSE is returned, the value is reversed to TRUE and **policy2** is evaluated.

```
export (!policy1 && policy2)
```

The sequential list [**policy1 policy2 policy3**] is not the same as the policy expression (**policy1 && policy2 && policy3**).

The sequential list is evaluated on the basis of a route matching a routing policy. For example, if **policy1** matches and the action is **accept** or **reject**, **policy2** and **policy3** are not evaluated. If **policy1** does not match, **policy2** is evaluated and so on until a match occurs and the action is **accept** or **reject**.

The policy expressions are evaluated on the basis of the action in a routing policy that is converted to the value of TRUE or FALSE and the logic of the specified logical operator. (For complete information about policy expression evaluation, see “How a Policy Expression Is Evaluated” on page 62.) For example, if **policy1** returns a value of FALSE, **policy2** and **policy3** are not evaluated. If **policy1** returns a value of TRUE, **policy2** is evaluated. If **policy2** returns a value of FALSE, **policy3** is not evaluated. If **policy2** returns a value of TRUE, **policy3** is evaluated.

You can also combine policy expressions and sequential lists. In the following example, if **policy1** returns a value of FALSE, **policy2** is evaluated. If **policy2** returns a value of TRUE and contains a **next policy** action, **policy3** is evaluated. If **policy2** returns a value of TRUE but does not contain an action, including a **next policy** action, **policy3** is still evaluated (because if you do not specify an action, **next term** or **next policy** are the default actions). If **policy2** returns a value of TRUE and contains an **accept** action, **policy3** is not evaluated.

```
export [(policy1 || policy2) policy3]
```

How a Policy Expression Is Evaluated

During evaluation, the policy framework software converts policy actions to values of TRUE or FALSE, which are factors in determining the flow control action that is performed upon a route. However, the software does not actually perform a flow control action on a route until it evaluates an entire policy expression.

The policy framework software evaluates a policy expression as follows:

1. The software evaluates a route against the first routing policy in a policy expression and converts the specified or default action to a value of TRUE or FALSE. (For information about the policy action conversion values, see Table 17 on page 60.)
2. The software takes the value of TRUE or FALSE and evaluates it against the logical operator used in the policy expression (see Table 18 on page 60). Based upon the logical operator used, the software determines whether or not to evaluate the next routing policy, if one is present.

The policy framework software uses a method of shortcut evaluation. When a result is certain, the software stops evaluating subsequent routing policies in the policy expression. For example, if the policy expression specifies logical AND and the evaluation of the first routing policy returns the value of FALSE, the software determines that the output will be FALSE no matter what the values of the unevaluated routing policies are. Therefore, the software does not evaluate the subsequent routing policies in this policy expression.

3. The software performs Step 1 and Step 2 for each subsequent routing policy in the policy expression, if they are present and if the software has determined that it is appropriate to evaluate them.
4. After evaluating the last routing policy, if it is appropriate, the software evaluates the value of TRUE or FALSE obtained from each routing policy evaluation. Based upon the logical operator used, it calculates an output of TRUE or FALSE.
5. The software converts the output of TRUE or FALSE back to an action. (For information about the policy action conversion values, see Table 17 on page 60.) The action is performed.

If each policy in the expression returned a value of TRUE, the software converts the output of TRUE back to the flow control action specified in the last policy. For example, if the policy expression (**policy1 && policy2**) is specified and **policy1** specifies **accept** and **policy2** specifies **next term**, the **next term** action is performed.

If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during

the evaluation of the remaining policies. For example, if the action specified in the first policy of a policy expression sets a route's metric to 500, this route matches the criteria of **metric 500** defined in the next policy. However, if a route characteristic manipulation action is specified in a policy located in the middle or the end of a policy expression, it is possible, because of the shortcut evaluation, that the policy is never evaluated and the manipulation of the route characteristic never occurs.

Example: Evaluating Policy Expressions

The following sample routing policy uses three policy expressions:

```
[edit]
policy-options {
  policy-statement policy-A {
    from {
      route-filter 10.10.0.0/16 orlonger;
    }
    then reject;
  }
}
policy-options {
  policy-statement policy-B {
    from {
      route-filter 10.20.0.0/16 orlonger;
    }
    then accept;
  }
}
protocols {
  bgp {
    neighbor 192.168.1.1 {
      export (policy-A && policy-B);
    }
    neighbor 192.168.2.1 {
      export (policy-A || policy-B);
    }
    neighbor 192.168.3.1 {
      export (!policy-A);
    }
  }
}
```

The policy framework software evaluates the transit BGP route **10.10.1.0/24** against the three policy expressions specified in the sample routing policy as follows:

- **(policy-A && policy-B)**—**10.10.1.0/24** is evaluated against **policy-A**. **10.10.1.0/24** matches the route list specified in **policy-A**, so the specified action of **reject** is returned. **reject** is converted to a value of **FALSE**, and **FALSE** is evaluated against the specified logical AND. Because the result of **FALSE** is certain no matter what the results of the evaluation of **policy-B** are (in policy expression logic, any result AND a value of **FALSE** produces the output of **FALSE**), **policy-B** is not evaluated and the output of **FALSE** is produced. The **FALSE** output is converted to **reject**, and **10.10.1.0/24** is rejected.

- (policy-A || policy-B)—10.10.1.0/24 is evaluated against policy-A. 10.10.1.0/24 matches the route list specified in policy-A, so the specified action of **reject** is returned. **reject** is converted to a value of FALSE, then FALSE is evaluated against the specified logical OR. Because logical OR requires at least one value of TRUE to produce an output of TRUE, 10.10.1.0/24 is evaluated against policy-B. 10.10.1.0/24 does not match policy-B, so the default action of **next-policy** is returned. The **next-policy** is converted to a value of TRUE, then the value of FALSE (for policy-A evaluation) and TRUE (for policy-B evaluation) are evaluated against the specified logical OR. In policy expression logic, FALSE OR TRUE produce an output of TRUE. The output of TRUE is converted to **next-policy**. (TRUE is converted to **next-policy** because **next-policy** was the last action retained by the policy framework software.) policy-B is the last routing policy in the policy expression, so the action specified by the default export policy for BGP, **accept**, is taken.
- (!policy-A)—10.10.1.0/24 is evaluated against policy-A. 10.10.1.0/24 matches the route list specified in policy-A, so the specified action of **reject** is returned. **reject** is converted to a value of FALSE, and FALSE is evaluated against the specified logical NOT. The value of FALSE is reversed to an output of TRUE based on the rules of logical NOT. The output of TRUE is converted to **accept**, and route 10.10.1.0/24 is accepted.

Side Effects of Omitting the "from" Statement from an Export Policy

In export policies, omitting the **from** statement in a term might lead to unexpected results. By default, if you omit the **from** statement, all routes are considered to match. For example, static and direct routes are not exported by BGP by default. However, if you create a term with an empty **from** statement, these routes inadvertently could be exported because they matched the **from** statement. For example, the following routing policy is designed to reject a few route ranges and then export routes learned by BGP (which is the default export behavior):

```
[edit]
routing-options {
  autonomous-system 56;
}
protocols {
  bgp {
    group 4 {
      export statics-policy;
      type external;
      peer-as 47;
      neighbor 192.168.1.1;
    }
  }
}
policy-options {
  policy-statement statics-policy {
    term term1 {
      from {
        route-filter 192.168.0.0/16 orlonger;
        route-filter 172.16.1.1/3 orlonger;
      }
      then reject; # reject the prefixes in the route list
    }
  }
}
```

```

    }
    term term2 {
      then {
        accept; # accept all other routes, including static and direct routes
      }
    }
  }
}

```

However, this routing policy results in BGP advertising static and direct routes to its peers because:

- **term1** rejects the destination prefixes enumerated in the route list.
- **term2**, because it has no **from** statement, matches all other routes, including static and direct routes, and accepts all these routes (with the **accept** statement).

To modify the preceding routing policy so that an IGP does not export unwanted routes, you can specify the following additional terms:

```

[edit]
routing-options {
  autonomous-system 56;
}
protocols {
  isis {
    export statics-policy;
  }
}
policy-options {
  policy-statement statics-policy {
    term term1 {
      from {
        route-filter 192.168.0.0/16 orlonger;
        route-filter 172.16.1.1/3 orlonger;
      }
      then reject; # reject the prefixes in the route list
    }
    term term2 { # reject direct routes
      from protocol direct;
      then reject;
    }
    term term3 { # reject static routes
      from protocol static;
      then reject;
    }
    term term4 { # reject local routes
      from protocol local;
      then reject;
    }
    term term5 { # reject aggregate routes
      from protocol aggregate;
      then reject;
    }
    term term6 {
      then accept; # accept all other routes
    }
  }
}

```

```

    }
  }
}

```

Applying Routing Policies to the Forwarding Table

To apply an export routing policy to the forwarding table, include the **export** statement:

```
export [ policy-names ];
```

You can include this statement at the following hierarchy levels:

- [edit routing-options forwarding-table]
- [edit logical-systems *logical-system-name* routing-options forwarding-table]

In the **export** statement, list the name of the routing policy to be evaluated when routes are being exported from the routing table into the forwarding table. Only active routes are exported from the routing table.

You can reference the same routing policy one or more times in the same or a different **export** statement.

For information about how the policy framework software evaluates a routing policy, see “How a Routing Policy Is Evaluated” on page 28.

You can apply export policies to routes being exported from the routing table into the forwarding table for the following features:

- Per-packet load balancing
- Class of service (CoS)

For more information about per-packet load balancing, see “Configuring the Load-Balance Per-Packet Action” on page 136. For more information about CoS, see the *JUNOS Class of Service Configuration Guide*.

Examples: Applying Routing Policies

Configure the router to export to IS-IS the routes that match the **dmz** and **local-customers** routing policies:

```

[edit]
protocols {
  isis {
    export [ dmz local-customers ];
  }
}

```

For three BGP peer groups, apply various export and import filters:

```

[edit]
protocols {
  bgp {

```



```

group 1 {
    type external;
    peer-as 47;
    export local-customers;
    import [ martian-filter long-prefix-filter as47-filter ];
    neighbor 192.168.1.4;
    neighbor 192.168.1.5;
}
group 2 {
    type external;
    peer-as 42;
    export local-customers;
    import [ martian-filter long-prefix-filter as42-filter ];
    neighbor 192.168.1.4;
    neighbor 192.168.1.5;
}
group 3 {
    type internal;
    export local-customers;
    neighbor 10.1.1.1;
}
}

```

Apply the long-prefix-filter prefix only to routes learned from a particular peer within a group:

```

[edit]
protocols {
    bgp {
        group 4 {
            type external;
            peer-as 47;
            export local-customers;
            import [ martian-filter as47-filter ];
            neighbor 192.168.1.4;
            neighbor 192.168.1.5;
            neighbor 192.168.1.6 {
                import [ martian-filter as47-filter long-prefix-filter ];
            }
        }
    }
}

```

Examples: Routing Policy Configuration

The following examples show how to configure routing policies for various purposes:

- Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS on page 68
- Example: Redistributing OSPF Routes into BGP on page 68
- Example: Exporting Direct Routes Into IS-IS on page 69
- Example: Exporting Internal IS-IS Level 1 Routes to Level 2 on page 69

- Example: Exporting IS-IS Level 2 Routes to Level 1 on page 70
- Example: Assigning Different Forwarding Next-Hop LSPs to Different Destination Prefixes on page 70
- Example: Grouping Destination Prefixes on page 71
- Example: Grouping Source Prefixes on page 72
- Example: Grouping Source and Destination Prefixes in a Forwarding Class on page 73
- Example: Accepting Routes with Specific Destination Prefixes on page 74
- Example: Accepting Routes from BGP with a Specific Destination Prefix on page 74

Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS

Redistribute BGP routes with a community tag of 444:5 into IS-IS, changing the metric to 14:

```
[edit]
protocols {
  isis {
    export edu-to-isis;
  }
}
policy-options {
  community edu members 444:5;
  policy-statement edu-to-isis {
    from {
      protocol bgp;
      community edu;
    }
    then {
      metric 14;
      accept;
    }
  }
}
```

Example: Redistributing OSPF Routes into BGP

Redistribute OSPF routes from Area 1 only into BGP, and do not advertise routes learned by BGP:

```
[edit]
routing-options {
  autonomous-system 56;
}
protocols {
  bgp {
    export ospf-into-bgp;
    group {
      type external;
      peer-as 23;
    }
  }
}
```

```

        allow {
            0.0.0.0/0;
        }
    }
}
}
policy-options {
    policy-statement ospf-into-bgp {
        term ospf-only {
            from {
                protocol ospf;
                area 1;
            }
            then accept;
        }
    }
}
}

```

Example: Exporting Direct Routes Into IS-IS

Export direct routes into IS-IS for all interfaces, even if IS-IS is not configured on an interface:

```

[edit]
protocols {
    isis {
        export direct-routes;
    }
}
policy-options {
    policy-statement direct-routes {
        from protocol direct;
        then accept;
    }
}
}

```

Example: Exporting Internal IS-IS Level 1 Routes to Level 2

Export IS-IS Level 1 internal-only routes into Level 2:

```

[edit]
protocols {
    isis {
        export L1-L2;
    }
}
policy-statement L1-L2 {
    term one {
        from {
            level 1;
            external;
        }
        then reject;
    }
    term two {

```

```

        from level 1;
        to level 2;
        then accept;
    }
}

```

Example: Exporting IS-IS Level 2 Routes to Level 1

Export IS-IS Level 2 routes into Level 1:

```

[edit]
protocols {
  isis {
    export L2-L1;
  }
}
policy-statement L2-L1 {
  term one {
    from level 2;
    to level 1;
    then accept;
  }
}

```

Example: Assigning Different Forwarding Next-Hop LSPs to Different Destination Prefixes

Assign different forwarding next-hop LSPs to different destination prefixes learned from BGP:

```

routing-options {
  router-id 10.10.20.101;
  autonomous-system 2;
  forwarding-table {
    export forwarding-policy;
  }
}
policy-options {
  policy-statement forwarding-policy {
    term one {
      from {
        protocol bgp;
        route-filter 10.1.0.0/16 orlonger;
      }
      then {
        install-nexthop lsp mc-c-lsp-1;
        accept;
      }
    }
    term two {
      from {
        protocol bgp;
        route-filter 10.2.0.0/16 orlonger;
      }
      then {
        install-nexthop lsp mc-c-lsp-2;
      }
    }
  }
}

```

```

        accept;
    }
}
term three {
    from {
        protocol bgp;
        route-filter 10.3.0.0/16 orlonger;
    }
    then {
        install-nexthop lsp mc-c-lsp-3;
        accept;
    }
}
}
}
}
protocols {
    mpls {
        label-switched-path mc-c-lsp-1 {
            from 10.10.20.101;
            to 10.10.20.103;
        }
        label-switched-path mc-c-lsp-2 {
            from 10.10.20.101;
            to 10.10.20.103;
        }
        label-switched-path mc-c-lsp-3 {
            from 10.10.20.101;
            to 10.10.20.103;
        }
    }
}
}

```

Example: Grouping Destination Prefixes

Configure a routing policy to group destination prefixes:

```

[edit]
policy-options {
    policy-statement set-dest-class {
        term 1 {
            from community nets1;
            then {
                destination-class on-net;
                accept;
            }
        }
        term 2 {
            from community nets2;
            then {
                destination-class off-net;
                accept;
            }
        }
    }
}
community nets1 [7:8 9:10];

```

```

    community nets2 [1:2 4:5];
}

```

Apply a routing policy to the forwarding table with the corresponding destination class:

```

[edit]
routing-options {
  forwarding-table {
    export set-dest-class;
  }
}

```

Enable packet counting on an interface:

```

[edit interfaces]
interfaces so-1/0/1 {
  unit 0 {
    family inet6 {
      accounting {
        destination-class-usage;
      }
    }
  }
}

```

Example: Grouping Source Prefixes

Configure a routing policy to group source prefixes, and allow prefixes that match the policy statement to have a source class created for them:

```

[edit]
policy-options {
  policy-statement set-gold-class {
    term {
      from
        route-filter 10.210.0.0/16 orlonger;
        route-filter 10.215.0.0/16 orlonger;
      then {
        source-class gold-class;
      }
    }
  }
}

```

Apply a routing policy to the forwarding table with the corresponding source class:

```

[edit]
routing-options {
  forwarding-table {
    export set-gold-class;
  }
}

```

Enable packet counting on an interface. In this example, one interface accommodates both input and output:

```
[edit interfaces]
interfaces ge/0/0/0 {
  unit 0 {
    family inet {
      accounting {
        source-class-usage {
          input;
          output;
        }
      }
    }
  }
}
```

Example: Grouping Source and Destination Prefixes in a Forwarding Class

Configure a routing policy to group source and destination prefixes in a forwarding class:

```
[edit]
policy-options {
  policy-statement set-bronze-class {
    term {
      from
      route-filter 10.210.0.0/16 orlonger;
      route-filter 10.215.0.0/16 orlonger;
      then {
        forwarding-class bronze-class;
      }
    }
  }
}
```

Apply a routing policy to the forwarding table with the corresponding forwarding class:

```
[edit]
routing-options {
  forwarding-table {
    export set-bronze-class;
  }
}
```

Enable counting of incoming source packets on an interface:

```
[edit interfaces]
interfaces fe/1/0/0 {
  unit 0 {
    family inet {
      accounting {
        source-class-usage {
          input;
        }
      }
    }
  }
}
```

```

    }
}

}

}

interfaces fe/1/0/1 {
    unit 0 {
        family inet {
            accounting {
                source-class-usage {
                    output;
                }
            }
        }
    }
}

}

}

}

interfaces fe/1/0/2 {
    unit 0 {
        family inet {
            accounting {
                destination-class-usage;
            }
        }
    }
}

}
```

Example: Accepting Routes with Specific Destination Prefixes

Accept routes with destination prefixes 201:db8::8000/32 and 201:db8::8001/32:

```
[edit policy-options]
policy-statement export_exact {
  term a {
    from {
      route-filter 201:db8::8000/32 exact;
      route-filter 201:db8::8001/32 exact;
    }
    then {
      accept;
    }
  }
  term b {
    then {
      reject;
    }
  }
}
```

Example: Accepting Routes from BGP with a Specific Destination Prefix

Accepts routes from BGP that have destination prefix 201:db8::8000/32:

```
[edit policy-options]
policy-statement export_exact {
  term a {
```




```

    from {
        protocol bgp;
        route-filter 201:db8::8000/32 exact;
    }
    then {
        accept;
    }
}
term b {
    then {
        reject;
    }
}
}

```

Example: ISP Network Case Study

This section provides an example of how policies might be used in a typical Internet service provider (ISP) network. In this network example (see Figure 11 on page 76), the ISP's AS number is 1000. The ISP has two transit peers (AS 11111 and AS 22222) to which it connects at an exchange point. The ISP is also connected to two private peers (AS 7000 and AS 8000) with which it exchanges specific customer routes. The ISP has two customers (AS 1234 and AS 2468) to which it connects using the BGP protocol.

Figure 11: ISP Network Example

In this example, the ISP policies are configured in an outbound direction; that is, the example focuses on the routes that the ISP announces to its peers and customers, and includes the following:

1. The ISP has been assigned AS 1000 and the routing space of 192.168.0/17. With the exception of the two customer networks shown in Figure 11 on page 76, all other customer routes are simulated with static routes.
2. The ISP has connectivity to two different exchange peers: AS 11111 and AS 22222. These peers are used for transit service to other portions of the Internet. This means that the ISP is accepting all routes (the full Internet routing table) from those BGP peers. To help maintain an optimized Internet routing table, the ISP is configured to advertise only two aggregate routes to the transit peers.
3. The ISP also has direct connectivity to two private peers: AS 7000 and AS 8000. The ISP administrators want all data to the private peers to use this direct link. As a result, all the customer routes from the ISP are advertised to those private peers. These peers then advertise all their customer routes to the ISP.
4. Finally, the ISP has two customers with which it communicates using BGP: AS 1234 and AS 2468. Each customer has a different set of requirements.

The following sections discuss the following topics:

- Requesting a Single Default Route on the Customer 1 Router on page 77
- Requesting Specific Routes on the Customer 2 Router on page 78
- Configuring a Peer Policy on ISP Router 3 on page 80
- Configuring Private and Exchange Peers on ISP Router 1 and 2 on page 82
- Configuring Locally Defined Static Routes on the Exchange Peer 2 Router on page 85
- Configuring Outbound and Generated Routes on the Private Peer 2 Router on page 85

Requesting a Single Default Route on the Customer 1 Router

Customer 1 has only a single route to the ISP and is using the ISP for transit service. This customer has requested a single default route (0.0.0.0/0) from the ISP.

```
[edit]
interfaces {
  so-0/0/1 {
    description "Connection to ISP Router 3";
    unit 0 {
      family inet {
        address 10.222.70.1/30;
      }
    }
  }
  fxp0 {
    description "MGMT INTERFACE - DO NOT DELETE";
    unit 0 {
      family inet {
        address 10.251.0.9/24;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.16.1/32;
      }
    }
  }
}
routing-options {
  static {
    route 192.168.16.0/27 reject;
    route 192.168.16.32/27 reject;
    route 192.168.16.64/27 reject;
    route 192.168.16.96/27 reject;
    route 192.168.16.128/27 reject;
    route 192.168.16.160/27 reject;
    route 192.168.16.192/27 reject;
  }
}
autonomous-system 1234;
```

```

}
protocols {
  bgp {
    group AS1000-Peers {
      type external;
      export send-statics;
      peer-as 1000;
      neighbor 10.222.70.2;
    }
  }
}
policy-options {
  policy-statement send-statics {
    term static-routes {
      from protocol static;
      then accept;
    }
  }
}

```

Requesting Specific Routes on the Customer 2 Router

Customer 2 has a link to the ISP, as well as a link to AS 8000. This customer has requested specific customer routes from the ISP, as well as from AS 8000. Customer 2 wants to use the ISP for transit service to the Internet, and has requested a default route from the ISP.

```

[edit]
interfaces {
  so-0/0/1 {
    description "Connection to ISP Router 3";
    unit 0 {
      family inet {
        address 10.222.61.2/30;
      }
    }
  }
  so-0/0/2 {
    description "Connection to Private-Peer 2";
    unit 0 {
      family inet {
        address 10.222.6.1/30;
      }
    }
  }
  fxp0 {
    description "MGMT INTERFACE - DO NOT DELETE";
    unit 0 {
      family inet {
        address 10.251.0.8/24;
      }
    }
  }
  lo0 {
    unit 0 {

```

```

        family inet {
            address 192.168.64.1/32;
        }
    }
}
routing-options {
    static {
        route 192.168.64.0/25 reject;
        route 192.168.64.128/25 reject;
        route 192.168.65.0/25 reject;
        route 192.168.66.0/25 reject;
        route 192.168.67.0/25 reject;
        route 192.168.65.128/25 reject;
        route 192.168.66.128/25 reject;
        route 192.168.67.128/25 reject;
    }
    autonomous-system 2468;
}
protocols {
    bgp {
        group External-Peers {
            type external;
            import inbound-routes;
            export outbound-routes;
            neighbor 10.222.61.1 {
                peer-as 1000;
            }
            neighbor 10.222.6.2 {
                peer-as 8000;
            }
        }
    }
}
policy-options {
    policy-statement outbound-routes {
        term statics {
            from protocol static;
            then accept;
        }
        term internal-bgp-routes {
            from {
                protocol bgp;
                as-path my-own-routes;
            }
            then accept;
        }
        term no-transit {
            then reject;
        }
    }
    policy-statement inbound-routes {
        term AS1000-primary {
            from {
                protocol bgp;
                as-path AS1000-routes;
            }
        }
    }
}

```

```

    }
    then {
        local-preference 200;
        accept;
    }
}
term AS8000-backup {
    from {
        protocol bgp;
        as-path AS8000-routes;
    }
    then {
        local-preference 50;
        accept;
    }
}
}
as-path my-own-routes "()";
as-path AS1000-routes "1000 .*";
as-path AS8000-routes "8000 .*";
}

```

Configuring a Peer Policy on ISP Router 3

On ISP Router 3, a separate policy is in place for each customer. The default route for Customer 1 is being sent by the **customer-1-peer** policy. This policy finds the 0.0.0.0/0 default route in **inet.0** and accepts it. The policy also rejects all other routes, thereby not sending all BGP routes on the ISP router. The **customer-2-peer** policy is for Customer 2 and contains the same policy terms, which also send the default route and no other transit BGP routes. The additional terms in the **customer-2-peer** policy send the ISP customer routes to Customer 2. Because there are local static routes on ISP router 3 that represent local customers, these routes are sent as well as all other internal (192.168.0/17) routes announced to the local router by the other ISP routers.

```

[edit]
routing-options {
    static { # simulate local customer routes
        route 192.168.72.0/22 reject;
        route 192.168.76.0/22 reject;
        route 192.168.80.0/22 reject;
        route 192.168.84.0/22 reject;
        route 192.168.88.0/22 reject;
        route 192.168.92.0/22 reject;
        route 192.168.72.0/21 reject;
        route 192.168.80.0/21 reject;
        route 192.168.88.0/21 reject;
    }
    generate { # install a default route if certain routes
        route 0.0.0.0/0 policy if-upstream-routes-exist; # from the exchange peers are
        advertised using BGP
    }
    autonomous-system 1000;
}
protocols {
    bgp {

```

```

group Internal-Peers {
    type internal;
    local-address 192.168.0.3;
    export internal-peers;
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group Customer-2 {
    type external;
    export customer-2-peer;
    peer-as 2468;
    neighbor 10.222.61.2;
}
group Customer-1 {
    type external;
    export customer-1-peer;
    peer-as 1234;
    neighbor 10.222.70.1;
}
}
isis {
    level 1 disable;
    interface so-0/0/0.0;
    interface ge-0/1/0.0;
    interface lo0.0;
}
}
policy-options {
    policy-statement internal-peers { # advertise local customer routes to peers
        term statics {
            from protocol static;
            then accept;
        }
        term next hop self { # set the BGP routes next hop to self for EBGp
            then { # routes advertised to IBGP peers
                next-hop self;
            }
        }
    }
}
policy-statement if-upstream-routes-exist {
    term only-certain-contributing-routes {
        from { # allow either the 10.100.0.0/17 or the 10.101.0.0/27 route
            route-filter 10.100.0.0/17 exact; # route to activate the generated route
            route-filter 10.101.0.0/27 exact; # route to activate the generated route
        }
        then accept; # do not allow any other route to activate
    } # the generated route in the routing table
    term reject-all-other-routes {
        then reject;
    }
}
}
policy-statement customer-2-peer { # advertise customer routes to all peers
    term statics {
        from protocol static;
        then accept;
    }
}

```

```

term-isp-and-customer routes { # advertise internal AS 1000 customer
    from { # to the customer
        protocol-bgp;
        route-filter 192.168.0.0/17 orlonger;
    }
    then accept;
}# advertise just the default route to AS 2468
term default-route {
    from {
        route-filter 0.0.0.0/exact;
    }
    then accept;
}
term reject-all-other-routes { # do not advertise any other routes
    then reject;
}
}
policy-statement customer-1-peer {
    term default-route { # advertise just the default route to AS 1234
        from {
            route-filter 0.0.0.0/0 exact;
        }
        then accept;
    }
    term reject-all-other-routes { # do not advertise any other routes
        then reject;
    }
}
}

```

Configuring Private and Exchange Peers on ISP Router 1 and 2

ISP Router 1 and ISP Router 2 each have two policies configured: the **private-peers** policy and the **exchange-peers** policy. Because of their similar configurations, this example describes the configuration for only ISP Router 2.

On ISP Router 2, the **private-peers** policy sends the ISP customer routes to the Private Peer 2 router. The policy accepts all local static routes (local ISP Router 2 customers) and all BGP routes in the **192.168.0/17** range (advertised by other ISP routers). These two terms represent the ISP customer routes. The final term rejects all other routes, which includes the entire Internet routing table sent by the exchange peers. These routes do not need to be sent to Private Peer 2 for two reasons:

- The peer already maintains a connection to Exchange Peer 2 in our example, so the routes are redundant.
- The Private Peer wants customer routes only. The **private-peers** policy accomplishes this goal. The **exchange-peers** policy sends routes to the Exchange Peer 2 router.

In the example, only two routes need to be sent to Exchange Peer 2:

- The aggregate route that represents the AS 1000 routing space of **192.168.0/17**. This route is configured as an aggregate route locally and is advertised by the **exchange-peers** policy.

- The address space assigned to Customer 2, **192.168.64/22**. This smaller aggregate route needs to be sent to Exchange Peer 2 because the customer is also attached to the AS 8000 peer (Private Peer 2).

Sending these two routes to Exchange Peer 2 allows other networks in the Internet to reach the customer through either the ISP or the Private Peer. If just the Private Peer were to advertise the **/22** network while the ISP maintained only its **/17** aggregate, then all traffic destined for the customer would transit AS 8000 only. Because the customer also wants routes from the ISP, the **192.168.64/22** route is announced by ISP Router 2. Like the larger aggregate route, the **192.168.64/22** route is configured locally and is advertised by the **exchange-peers** policy. The final term in that policy rejects all routes, including the specific customer networks of the ISP, the customer routes from Private Peer 1, the customer routes from Private Peer 2, and the routing table from Exchange Peer 1. In essence, this final term prevents the ISP from performing transit services for the Internet at large.

```
[edit]
routing-options {
  static {
    route 192.168.32.0/22 reject;
    route 192.168.36.0/22 reject;
    route 192.168.40.0/22 reject;
    route 192.168.44.0/22 reject;
    route 192.168.48.0/22 reject;
    route 192.168.52.0/22 reject;
    route 192.168.32.0/21 reject;
    route 192.168.40.0/21 reject;
    route 192.168.48.0/21 reject;
  }
  aggregate {
    route 192.168.0.0/17;
    route 192.168.64.0/22;
  }
  autonomous-system 1000;
}
protocols {
  bgp {
    group Internal-Peers {
      type internal;
      local-address 192.168.0.2;
      export internal-peers;
      neighbor 192.168.0.1;
      neighbor 192.168.0.3;
    }
    group AS8000-Peers {
      type external;
      export private-peers;
      peer-as 8000;
      neighbor 10.222.45.2;
    }
    group AS22222-Peers {
      type external;
      export exchange-peers;
      peer-as 22222;
      neighbor 10.222.46.1;
```

```

    }
  }
  isis {
    level 1 disable;
    interface so-0/0/0.0;
    interface ge-0/2/0.0;
    interface lo0.0;
  }
}
policy-options {
  policy-statement internal-peers {
    term statics {
      from protocol static;
      then accept;
    }
    term next-hop-self {
      then {
        next-hop self;
      }
    }
  }
}
policy-statement private-peers {
  term statics {
    from protocol static;
    then accept;
  }
  term isp-and-customer-routes {
    from {
      protocol bgp;
      route-filter 192.168.0.0/17 orlonger;
    }
    then accept;
  }
  term reject-all {
    then reject;
  }
}
policy-statement exchange-peers {
  term AS1000-Aggregate {
    from {
      protocol aggregate;
      route-filter 192.168.0.0/17 exact;
    }
    then accept;
  }
  term Customer-2-Aggregate {
    from {
      protocol aggregate;
      route-filter 192.168.64.0/22 exact;
    }
    then accept;
  }
  term reject-all-other-routes {
    then reject;
  }
}

```

```
}
```

Configuring Locally Defined Static Routes on the Exchange Peer 2 Router

The Exchange Peer 2 router exchanges all routes with all BGP peers. The outbound-routes policy for Exchange Peer 2 advertises locally defined static routes using BGP.

```
[edit]
protocols {
  bgp {
    group Peers {
      type external;
      export outbound-routes;
      neighbor 10.222.4.1 {
        peer-as 11111;
      }
      neighbor 10.222.44.2 {
        peer-as 8000;
      }
      neighbor 10.222.46.2 {
        peer-as 1000;
      }
    }
  }
}
policy-options {
  policy-statement outbound-routes { # advertise the simulated Internet routes
    term statics { # to all BGP peers
      from protocol static;
      then accept;
    }
  }
}
```

Configuring Outbound and Generated Routes on the Private Peer 2 Router

The Private Peer 2 router performs two main functions:

- Advertises routes local to AS 8000 to both the Exchange Peers and the ISP routers. The outbound-routes policy advertises the local static routes (that is, customers) on the router, and also advertises all routes learned by BGP that originated in either AS 8000 or AS 2468. These routes include other AS 8000 customer routes in addition to the AS 2468 customer. The AS routes are identified by an AS path regular expression match criteria in the policy.
- Advertises the 0.0.0.0/0 default route to the AS 2468 customer router. To accomplish this, the Private Peer creates a generated route for 0.0.0.0/0 locally on the router. This generated route is further assigned a policy called *if-upstream-routes-exist*, which allows only certain routes to contribute to the generated route, making it an active route in the routing table. Once the route is active, it can be sent to the AS 2468 router using BGP and the configured policies. The *if-upstream-routes-exist* policy accepts only the 20.100.0.0/17 route from Exchange Peer 2, and rejects all other routes. If the 20.100.0.0/17 route

is withdrawn by the Exchange Peer, the Private Peer loses the 0.0.0.0/0 default route and withdraws the default route from the AS 2468 customer router.

```
[edit]
routing-options { # simulate local customer routes
  static {
    route 172.16.64.0/20 reject;
    route 172.16.80.0/20 reject;
    route 172.16.96.0/20 reject;
    route 172.16.112.0/20 reject;
    route 172.16.72.0/21 reject;
    route 172.16.88.0/21 reject;
    route 172.16.104.0/21 reject;
    route 172.16.120.0/21 reject;
  }
  generate {
    route 0.0.0.0/0 policy if-upstream-routes-exist;
  }
  autonomous-system 8000;
  protocols {
    bgp {
      group External-Peers {
        type external;
        export outbound-routes;
        neighbor 10.222.44.1 {
          peer-as 22222;
        }
        neighbor 10.222.45.1 {
          peer-as 1000;
        }
      }
      group Customers {
        type external;
        export internal-routes;
        neighbor 10.222.6.1 {
          peer-as 2468;
        }
      }
    }
  }
  policy-options {
    policy-statement outbound-routes { # advertise local customer routes
      term statics {
        from protocol static;
        then accept;
      }
      term allowed-bgp-routes {
        from { # advertise routes
          as-path [ my-own-routes AS2468-routes ];
        }
        then accept;
      }
      term no-transit {
        then reject; # do not advertise any other routes
      }
    }
  }
}
```

```

policy-statement internal-routes { # advertise local customer routes
    term statics {
        from protocol static;
        then accept;
    }
    term default-route { # advertise just the default route
        from {
            route-filter 0.0.0.0/0 exact;
        }
        then accept;
    }
    term reject-all-other-routes { # do not advertise any other routes
        then reject;
    }
}
policy-statement if-upstream-routes-exist {
    term as-22222-routes {
        from { # allow the 10.100.0.0/17 route to activate
            route-filter 10.100.0.0/17 exact; # the generated route in the routing
            # table
        }
        then accept;
    }
    term reject-all-other-routes {
        then reject; # do not allow any other route to activate
    } # the generated route in the routing table
}
as-path my-own-routes "()";
as-path AS2468-routes "2468";
}

```

Configuring the Discard Interface

The discard interface allows you to protect a network from denial-of-service (DoS) attacks by identifying the target IP address that is being attacked and configuring a policy to forward all packets to a discard interface. All packets forwarded to the discard interface are dropped.

To configure the discard interface, include the **dsc** statement:

```

dsc {
    unit 0 {
        family inet {
            filter {
                input filter-name;
                output filter-name;
            }
        }
    }
}

```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name*]

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

The **dsc** interface name denotes the discard interface. The discard interface supports only unit 0. For more information about configuring interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

The following two configurations are required to configure a policy to forward all packets to the discard interface.

Configure an input policy to associate a community with the discard interface:

```
[edit]
policy-options {
  community community-name members [ community-id ];
  policy-statement statement-name {
    term term-name {
      from community community-name;
      then {
        next-hop address; # Remote end of the point-to-point interface
        accept;
      }
    }
  }
}
```

Configure an output policy to set up the community on the routes injected into the network:

```
[edit]
policy-options {
  policy-statement statement-name {
    term term-name {
      from prefix-list name;
      then community (set | add | delete) community-name;
    }
  }
}
```

Testing Routing Policies

Before applying a routing policy in a live environment, you can use the **test policy** command to ensure that the policy produces the results that you expect. The **test policy** command uses the following syntax:

```
user@host> test policy policy-name prefix
```

For more information about test commands, see the *JUNOS Routing Protocols and Policies Command Reference*.

Example: Testing a Routing Policy

Test the following policy, which looks for unwanted routes and rejects them:

```
[edit policy-options]
policy-statement reject-unwanted-routes {
  term drop-these-routes {
    from {
      route-filter 0/0 exact;
      route-filter 10/8 orlonger;
      route-filter 172.16/12 orlonger;
      route-filter 192.168/16 orlonger;
      route-filter 224/3 orlonger;
    }
    then reject;
  }
}
```

Test this policy against all routes in the routing table:

```
user@host> test policy reject-unwanted-routes 0/0
```

Test this policy against a specific set of routes:

```
user@host> test policy reject-unwanted-routes 10.49.0.0/16
```


Chapter 5

Extended Match Conditions Configuration

This chapter describes how to configure extended match conditions for a routing policy. In general, the extended match conditions include criteria that are defined separately from the routing policy (autonomous system [AS] path regular expressions, communities, and prefix lists) and are more complex than standard match conditions. You should have an in-depth understanding of extended match conditions before configuring them. For information about standard match conditions, see Table 14 on page 42.

This chapter describes how to configure extended match conditions:

- Configuring AS Path Regular Expressions on page 91
- Configuring Communities on page 97
- Configuring Prefix Lists on page 110
- Configuring Route Lists on page 114
- Configuring Subroutines on page 123
- Configuring the Condition Statement on page 127

Configuring AS Path Regular Expressions

A Border Gateway Protocol (BGP) *AS path* is a path to a destination. An AS path consists of the AS numbers of networks that a packet traverses if it takes the associated route to a destination. The AS numbers are assembled in a sequence, or path, that is read from right to left. For example, for a packet to reach a destination using a route with an AS path **5 4 3 2 1**, the packet first traverses AS 1 and so on until it reaches AS 5, which is the last AS before its destination.

You can define a match condition based on all or portions of the AS path. To do this, you create a named AS path regular expression and then include it in a routing policy.

The following sections discuss the following tasks for configuring AS path regular expressions and provides the following examples:

- Defining AS Path Regular Expressions on page 92
- How AS Path Regular Expressions Are Evaluated on page 96
- Examples: Configuring AS Path Regular Expressions on page 96

Defining AS Path Regular Expressions

You can create a named AS path regular expression and then include it in a routing policy with the **as-path** match condition (described in Table 14 on page 42). To create a named AS path regular expression, include the **as-path** statement:

```
as-path name regular-expression;
```

You can include this statement at the following hierarchy levels:

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

To include the AS path regular expression in a routing policy, include the **as-path** match condition in the **from** statement:

```
as-path name regular-expression;
policy-statement policy-name {
  term term-name {
    from {
      names;
    }
  }
}
```

Additionally, you can create a named AS path group made up of AS path regular expressions and then include it in a routing policy with the **as-path-group** match condition. To create a named AS path group, include the **as-path-group** statement:

```
as-path-group group-name {
  name [ regular-expressions ];
}
```

You can include this statement at the following hierarchy levels:

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

To include the AS path regular expressions within the AS path group in a routing policy, include the **as-path-group** match condition in the **from** statement:

```
as-path-group group-name {
  name [ regular-expressions ];
}
policy-statement policy-name {
  term term-name {
    from {
      as-path-group group-name;
    }
  }
}
```



NOTE: You cannot have both **as-path** and **as-path-group** in the same policy term.



NOTE: You can include the names of multiple AS path regular expressions in the **as-path** match condition in the **from** statement. If you do this, only one AS path regular expression needs to match for a match to occur. The AS path regular expression matching is effectively a logical OR operation.

The AS path name identifies the regular expression. It can contain letters, numbers, and hyphens (-), and can be up to 255 characters. To include spaces in the name, enclose the entire name in quotation marks (" ").

The regular expression is used to match all or portions of the AS path. It consists of two components, which you specify in the following format:

term <operator>

- **term**—Identifies an AS. You can specify it in one of the following ways:
 - AS number—The entire AS number composes one term. You cannot reference individual characters within an AS number, which differs from regular expressions as defined in POSIX 1003.2.
 - Wildcard character—Matches any single AS number. The wildcard character is a period (.). You can specify multiple wildcard characters.
 - AS path—A single AS number or a group of AS numbers enclosed in parentheses. Grouping the regular expression in this way allows you to perform a common operation on the group as a whole and to give the group precedence. The grouped path can itself include operators.

Beginning with JUNOS Release 9.1, the range for AS numbers has been extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. You can now configure a value that is from 1 through 4,294,967,295. The JUNOS software continues to support 2-byte AS numbers. For more information about configuring the AS number for the local router and the Border Gateway Protocol (BGP), see the *JUNOS Routing Protocols Configuration Guide*.

- **operator**—(Optional) An operator specifying how the term must match. Most operators describe how many times the term must be found to be considered a match (for example, any number of occurrences, or zero, or one occurrence). Table 19 on page 94 lists the regular expression operators supported for AS paths. You place operators immediately after **term** with no intervening space, except for the pipe (|) and dash (-) operators, which you place between two terms, and parentheses, with which you enclose terms.

You can specify one or more term-operator pairs in a single regular expression.

Table 20 on page 94 shows examples of how to define regular expressions to match AS paths.

Table 19: AS Path Regular Expression Operators

Operator	Match Definition
$\{m,n\}$	At least m and at most n repetitions of <i>term</i> . Both m and n must be positive integers, and m must be smaller than n .
$\{m\}$	Exactly m repetitions of <i>term</i> . m must be a positive integer.
$\{m,\}$	m or more repetitions of <i>term</i> . m must be a positive integer.
$*$	Zero or more repetitions of <i>term</i> . This is equivalent to $\{0,\}$.
$+$	One or more repetitions of <i>term</i> . This is equivalent to $\{1,\}$.
$?$	Zero or one repetition of <i>term</i> . This is equivalent to $\{0,1\}$.
$ $	One of two terms on either side of the pipe.
$-$	Between a starting and ending range, inclusive.
$^$	A character at the beginning of a community attribute regular expression. This character is added implicitly; therefore, the use of it is optional.
$\$$	A character at the end of a community attribute regular expression. This character is added implicitly; therefore, the use of it is optional.
$()$	A group of terms that are enclosed in the parentheses. If enclosed in quotation marks with no intervening space ("0"), indicates a null. Intervening space between the parentheses and the terms is ignored.
$[]$	Set of characters. One character from the set can match. To specify the start and end of a range, use a hyphen (-).
$!$	The not operator.

Table 20: Examples of Defining AS Path Regular Expressions

AS Path to Match	Regular Expression	Sample Matches
AS path is 1234	1234	1234
Zero or more occurrences of AS number 1234	1234*	1234 1234 1234 1234 1234 1234 Null AS path
Zero or one occurrence of AS number 1234	1234? or 1234{0,1}	1234 Null AS path
One through four occurrences of AS number 1234	1234{1,4}	1234 1234 1234 1234 1234 1234 1234 1234 1234 1234

Table 20: Examples of Defining AS Path Regular Expressions *(continued)*

AS Path to Match	Regular Expression	Sample Matches
One through four occurrences of AS number 12, followed by one occurrence of AS number 34	12{1,4} 34	12 34 12 12 34 12 12 12 34 12 12 12 12 34
Range of AS numbers to match a single AS number	123–125 [123–125]*	123 or 124 or 125 Null AS path 123 124 124 125 125 125
Path whose second AS number must be 56 or 78	(. 56) (. 78) or . (56 78)	1234 56 34 78
Path whose second AS number might be 56 or 78	. (56 78)?	1234 56 34
Path whose first AS number is 123 and second AS number is either 56 or 78	123 (56 78)	123 56 123 78
Path of any length, except nonexistent, whose second AS number can be anything, including nonexistent	. * or .{0,}	1234 1234 5678 1234 5 6 7 8
AS path is 1 2 3	1 2 3	1 2 3
One occurrence of the AS numbers 1 and 2, followed by one or more occurrences of the number 3	1 2 3 +	1 2 3 1 2 3 3 1 2 3 3 3
One or more occurrences of AS number 1, followed by one or more occurrences of AS number 2, followed by one or more occurrences of AS number 3	1 + 2 + 3 +	1 2 3 1 1 2 3 1 1 2 2 3 1 1 2 2 3 3
Path of any length that begins with AS numbers 4, 5, 6	4 5 6 . *	4 5 6 4 5 6 7 8 9
Path of any length that ends with AS numbers 4, 5, 6	. * 4 5 6	4 5 6 1 2 3 4 5 6
AS path 5, 12, or 18	5 12 18	5 12 18

Null AS Path

You can use AS path regular expressions to create a null AS path that matches routes (prefixes) that have originated in your AS. These routes have not been advertised to your AS by any external peers. To create a null AS path, use the parentheses operator enclosed in quotation marks with no intervening spaces:

"()"

Example: Null AS Path

AS 1 and AS 3 are connected to AS 2, which you administrate. AS 3 advertises its routes to your AS, but you do not want to advertise AS 3 routes to AS 1 and thereby begin routing traffic from AS 1 to AS 3 through your AS. To prevent this situation from occurring, you can configure an export policy for AS 1 (1.2.2.6) that allows routes for your AS to be advertised to AS 1 but does not allow routes for AS 3 or routes for any other connected AS to be advertised to AS 1:

```
[edit policy-options]
null-as "()";
policy-statement only-my-routes {
  term just-my-as {
    from {
      protocol bgp;
      as-path null-as;
    }
    then accept;
  }
  term nothing-else {
    then reject;
  }
}
protocol {
  bgp {
    neighbor 10.2.2.6 {
      export only-my-routes;
    }
  }
}
```

How AS Path Regular Expressions Are Evaluated

AS path regular expressions implement the extended (modern) regular expressions as defined in POSIX 1003.2. They are identical to the UNIX regular expressions with the following exceptions:

- The basic unit of matching in an AS path regular expression is the AS number and not an individual character.
- A regular expression matches a route only if the AS path in the route exactly matches *regular-expression*. The equivalent UNIX regular expression is *^regular-expression\$*. For example, the AS path regular expression **1234** is equivalent to the UNIX regular expression **^1234\$**.
- You can specify a regular expression using wildcard operators.

Examples: Configuring AS Path Regular Expressions

Exactly match routes with the AS path **1234 56 78 9** and accept them:

```
[edit]
```

```

policy-options {
  wellington "1234 56 78 9";
  policy-statement from-wellington {
    term term1 {
      from as-path wellington;
    }
    then {
      preference 200;
      accept;
    }
    term term2 {
      then reject;
    }
  }
}

```

Match alternate paths to an AS and accept them after modifying the preference:

```

[edit]
policy-options {
  wellington-alternate "1234{1,6} (56|47)? (78|101|112)* 9+";
  policy-statement from-wellington {
    from as-path wellington-alternate;
  }
  then {
    preference 200;
    accept;
  }
}

```

Match routes with an AS path of 123, 124, or 125 and accept them after modifying the preference:

```

[edit]
policy-options {
  addison "123-125";
  policy-statement from-addison {
    from as-path addison;
  }
  then {
    preference 200;
    accept;
  }
}

```

Configuring Communities

A *BGP community* is a group of destinations that share a common property. Community information is included as a path attribute in BGP update messages. This information identifies community members and allows you to perform actions on a group without having to elaborate upon each member. You can define match conditions based on a BGP community, which this section describes. For information about actions that can be performed on the community, see Table 16 on page 49.



NOTE: You can assign community tags to non-BGP routes through configuration (for static, aggregate, or generated routes) or an import routing policy. These tags can then be matched when BGP exports the routes.

This section includes the following information:

- Defining Communities on page 98
- How Communities Are Evaluated on page 109

Defining Communities

You can create a named community and include it in a routing policy with the **community** match condition (described in Table 14 on page 42).

You can configure community and extended communities attributes to be included in BGP update messages. The community attribute is only four octets. The BGP extended communities attribute provides a larger range (eight octets) for grouping or categorizing communities. You can use community and extended communities attributes to trigger routing decisions, such as acceptance, rejection, preference, or redistribution.

The *community-ids* format varies according to the type of attribute that you use. The BGP community attribute format is *as-number:community-value*. The BGP extended communities attribute format is *type:administrator:assigned-number*.

When specifying *community-ids* for the community attribute, you can use UNIX-style regular expressions. Regular expressions are not supported for the extended communities attributes.

To define communities, you perform the following tasks:

- Configuring the Community Attribute on page 98
- Configuring the Extended Communities Attribute on page 107
- Inverting Community Matches on page 109
- Configuring Link Bandwidth on page 109

Configuring the Community Attribute

To create a named community and define the community members, include the **community** statement:

```
community name {
  invert-match;
  members [ community-ids ];
}
```

You can include this statement at the following hierarchy levels:

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

To include the community in a routing policy, include the **community** condition in the **from** statement:

```
[edit policy-options]
community name members [ community-ids ];
policy-statement policy-name {
  term term-name {
    from {
      community [ names ];
    }
  }
}
```



NOTE: You can include the names of multiple communities in the **community** match condition in the **from** statement. If you do this, only one community needs to match for a match to occur. The community matching is effectively a logical OR operation.

name identifies the community or communities. It can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").

community-ids defines one or more members of the community. It consists of two components, which you specify in the following format:

as-number:community-value;

- **as-number**—AS number of the community member. It can be a value from 0 through 65,535. For more information about configuring the AS number for the router and the Border Gateway Protocol (BGP), see the *JUNOS Routing Protocols Configuration Guide*. You can specify the AS number in one of the following ways:
 - AS number.
 - Asterisk (*)—A wildcard character that matches all AS numbers. (In the definition of the community attribute, the asterisk also functions as described in Table 21 on page 101.)
 - Period (.)—A wildcard character that matches any single digit in an AS number.
 - Group of AS numbers—A single AS number or a group of AS numbers enclosed in parentheses. Grouping the numbers in this way allows you to perform a common operation on the group as a whole and to give the group precedence. The grouped numbers can themselves include regular expression operators. For more information about the community regular expressions, see “Configuring the Community Attribute Using UNIX Regular Expressions” on page 100.
- **community-value**—Identifier of the community member. It can be a number from 0 through 65,535. You can specify the community value in one of the following ways:

- Community value number.
- Asterisk (*)—A wildcard character that matches all community values. (In the definition of the community attribute, the asterisk also functions as described in Table 21 on page 101.)
- Period (.)—A wildcard character that matches any single digit in a community value number.
- Group of community value numbers—A single community value number or a group of community value numbers enclosed in parentheses. Grouping the regular expression in this way allows you to perform a common operation on the group as a whole and to give the group precedence. The grouped path can itself include regular expression operators.

You also can specify *community-id* as one of the following well-known community names, which are defined in RFC 1997, *BGP Communities Attribute*:

- **no-advertise**—Routes in this community name must not be advertised to other BGP peers.
- **no-export**—Routes in this community must not be advertised outside a BGP confederation boundary.
- **no-export-subconfed**—Routes in this community must not be advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation.

Additionally, you can explicitly exclude BGP community information with a static route by using the **none** option. Include this option when configuring an individual route in the **route** portion to override a community option specified in the **defaults** portion.

See also “Configuring the Extended Communities Attribute” on page 107.

The following sections discuss the following topics:

- Configuring the Community Attribute Using UNIX Regular Expressions on page 100
- Do Not Advertise Communities to Neighbors on page 102
- Examples: Configuring the Community Attribute on page 103

Configuring the Community Attribute Using UNIX Regular Expressions

When specifying *community-ids*, you can use UNIX-style regular expressions to specify the AS number and the member identifier. A regular expression consists of two components, which you specify in the following format:

term operator;

term identifies the string to match.

operator specifies how the term must match. Table 21 on page 101 lists the regular expression operators supported for the community attribute. You place an operator immediately after *term* with no intervening space, except for the pipe (|) and dash (–) operators, which you place between two terms, and parentheses, with which you

enclose terms. Table 22 on page 102 shows examples of how to define *community-ids* using community regular expressions. The operator is optional.

Community regular expressions are identical to the UNIX regular expressions. Both implement the extended (or modern) regular expressions as defined in POSIX 1003.2.

Community regular expressions evaluate the string specified in *term* on a character-by-character basis. For example, if you specify **1234:5678** as *term*, the regular expressions see nine discrete characters, including the colon (:), instead of two sets of numbers (**1234** and **5678**) separated by a colon.



NOTE: Beginning with JUNOS Release 9.1, the range that you can configure for the AS number has been extended to provide BGP support for 4-byte AS numbers. For more information about configuring the AS number, see the *JUNOS Routing Protocols Configuration Guide*.

Table 21: Community Attribute Regular Expression Operators

Operator	Match Definition
$\{m,n\}$	At least <i>m</i> and at most <i>n</i> repetitions of <i>term</i> . Both <i>m</i> and <i>n</i> must be positive integers, and <i>m</i> must be smaller than <i>n</i> .
$\{m\}$	Exactly <i>m</i> repetitions of <i>term</i> . <i>m</i> must be a positive integer.
$\{m,\}$	<i>m</i> or more repetitions of <i>term</i> . <i>m</i> must be a positive integer.
*	Zero or more repetitions of <i>term</i> . This is equivalent to $\{0,\}$.
+	One or more repetitions of <i>term</i> . This is equivalent to $\{1,\}$.
?	Zero or one repetition of <i>term</i> . This is equivalent to $\{0,1\}$.
	One of the two terms on either side of the pipe.
–	Between a starting and ending range, inclusive.
^	Character at the beginning of a community attribute regular expression. We recommend the use of this operator for the clearest interpretation of your community attribute regular expression. If you do not use this operator, the regular expression 123:456 could also match a route tagged with 5123:456 .
\$	Character at the end of a community attribute regular expression. We recommend the use of this operator for the clearest interpretation of your community attribute regular expression. If you do not use this operator, the regular expression 123:456 could also match a route tagged with 123:4563 .
[]	Set of characters. One character from the set can match. To specify the start and end of a range, use a hyphen (-). To specify a set of characters that do not match, use the caret (^) as the first character after the opening square bracket ([).
()	Group of terms that are enclosed in parentheses. If enclosed in quotation marks with no intervening space ("()"), indicates a null. Intervening space between the parentheses and the terms is ignored.

Table 21: Community Attribute Regular Expression Operators (*continued*)

Operator	Match Definition
" "	Characters (such as space, tab, question mark, and bracket) that are enclosed within quotation marks in a community attribute regular expression indicate special characters.

Table 22: Examples of Defining Community Attribute Regular Expressions

Community Attribute to Match	Regular Expression	Sample Matches
AS number is 56 or 78. Community value is any number.	^((56) (78)):(.*)\$	56:1000 78:65000
AS number is 56. Community value is any number that starts with 2.	^56:(2.*)\$	56:2 56:222 56:234
AS number is any number. Community value is any number that ends with 5, 7, or 9.	^(.*):(.*[579])\$	1234:5 78:2357 34:65009
AS number is 56 or 78. Community value is any number that starts with 2 and ends with 2 through 8.	^((56) (78)):(2.*[2-8])\$	56:22 56:21197 78:2678

Do Not Advertise Communities to Neighbors

By default, communities are sent to BGP peers. To suppress the advertisement of communities to a neighbor, remove all communities. When the result of an export policy is an empty set of communities, the community attribute is not sent. To remove all communities, first define a wildcard set of communities (here, the community is named *wild*):

```
[edit policy-options]
community wild members "*" : "*";
```

Then, in the routing policy statement, specify the **community delete** action:

```
[edit policy-options]
policy-statement policy-name {
  term term-name {
    then community delete wild;
  }
}
```

To suppress a particular community from any AS, define the community as community wild members "*:community-value".

Examples: Configuring the Community Attribute

Create a community named dunedin and apply it in a routing policy statement:

```
[edit]
policy-options {
  community dunedin members [ 56:2379 23:46944 ];
  policy-statement from-dunedin {
    from community dunedin;
    then {
      metric 2;
      preference 100;
      next policy;
    }
  }
}
```

The preceding example modifies the metric and preference for routes that contain members of community **dunedin** only.



NOTE: You cannot set or add a community in a policy whose members use regular expressions or a wildcard.

Delete a particular community from a route, leaving remaining communities untouched:

```
[edit]
policy-options {
  community dunedin members 701:555;
  policy-statement delete-dunedin {
    then {
      community delete dunedin;
    }
  }
}
```

Remove any community from a route with the AS number of 65534 or 65535:

```
[edit]
policy-options {
  community my-as1-transit members [ 65535:10 65535:11 ];
  community my-as2-transit members [ 65534:10 65534:11 ];
  community my-wild members [ 65534:* 65535:* ];
  policy-statement delete-communities {
    from {
      community [ my-as1-transit my-as2-transit ];
    }
    then {
      community delete my-wild;
    }
  }
}
```

```
    }
  }
```

Match the set of community members 5000, 5010, 5020, 5030, and so on up to 5090:

```
[edit]
policy-options {
  community customers members "^1111:50.0$";
  policy-statement advertise-customers {
    from community customers;
    then accept;
  }
}
```

Reject routes that are longer than /19 in Class A space, /16 in Class B space, and /24 in Class C space:

```
[edit policy-options]
community auckland-accept members 555:1;
policy-statement drop-specific-routes {
  from {
    route-filter 0.0.0.0/1 upto /19 {
      community add auckland-accept;
      next policy;
    }
    route-filter 172.16.0.0/2 upto /16 {
      community add auckland-accept;
      next policy;
    }
    route-filter 192.168.0.0/3 upto /24 {
      community add auckland-accept;
      next policy;
    }
  }
  then reject;
}
```

In the preceding example, for routes that are not rejected, the tag **auckland-accept** is added.

Create routing policies to handle peer and customer communities. This example does the following:

- Customer routes that match the attributes defined in the **lcl20x-low** communities, for example, **lcl201-low**, are accepted and their local preference is changed to 80.
- Customer routes that match the attributes defined in the **lcl20x-high** communities, for example, **lcl201-high**, are accepted and have their local preference changed to 120.
- Internal routes that match the attributes defined in the **internal20x** communities, for example, **internal201**, are rejected and not advertised to customers.
- Routes received from a peer are assigned a metric of 10 and the community defined in **peer201**.

- Routes that match the attributes defined in the `prepend20x-x` communities, for example, `prepend201-1`, `prepend201-2`, or `prepend201-3`, are sent to peers and have the AS number 201 prepended the specified number of times.
- Routes that match the attributes defined in the `peer20x`, `custpeer20x`, and `internal20x` communities, for example, `peer201`, `custpeer201`, or `internal201`, respectively, are rejected and not advertised to peers.

[edit]

```

policy-options {
  community internal201 members 201:112;
  community internal202 members 202:112;
  community internal203 members 203:112;
  community internal204 members 204:112;
  community internal205 members 205:112;
  community peer201 members 201:555;
  community peer202 members 202:555;
  community peer203 members 203:555;
  community peer204 members 204:555;
  community peer205 members 205:555;
  community custpeer201 members 201:20;
  community custpeer202 members 202:20;
  community custpeer203 members 203:20;
  community custpeer204 members 204:20;
  community custpeer205 members 205:20;
  community prepend201-1 members 201:1;
  community prepend202-1 members 202:1;
  community prepend203-1 members 203:1;
  community prepend204-1 members 204:1;
  community prepend205-1 members 205:1;
  community prepend201-2 members 201:2;
  community prepend202-2 members 202:2;
  community prepend203-2 members 203:2;
  community prepend204-2 members 204:2;
  community prepend205-2 members 205:2;
  community prepend201-3 members 201:3;
  community prepend202-3 members 202:3;
  community prepend203-3 members 203:3;
  community prepend204-3 members 204:3;
  community prepend205-3 members 205:3;
  community lcl201-low members 201:80;
  community lcl202-low members 202:80;
  community lcl203-low members 203:80;
  community lcl204-low members 204:80;
  community lcl205-low members 205:80;
  community lcl 20x-high members "^20 [ 1-5 ] : 120$";
  policy-statement in-customer {
    term term1 {
      from {
        protocol bgp;
        community lcl 20x-high;
      }
      then {
        local-preference 80;
        accept;
      }
    }
  }
}

```

```

    }
    term term2 {
        from {
            protocol bgp;
            community [ lcl201-high lcl202-high lcl203-high lcl204-high lcl205-high
                ];
        }
        then local-preference 120;
    }
    then next policy;
}
policy-statement out-customer {
    term term1 {
        from {
            protocol bgp;
            community [internal201 internal202 internal203 internal204 internal205];
        }
        then reject;
    }
    then next policy;
}
policy-statement in-peer {
    from protocol bgp;
    then {
        metric 10;
        community set peer201;
    }
}
policy-statement out-peer {
    term term1 {
        from {
            protocol bgp;
            community [ prepend201-1 prepend202-1 prepend203-1 prepend204-1
                prepend205-1 ];
        }
        then as-path-prepend 201;
    }
    term term2 {
        from {
            protocol bgp;
            community [ prepend201-2 prepend202-2 prepend203-2 prepend204-2
                prepend205-2 ];
        }
        then as-path-prepend "201 201";
    }
    term term3 {
        from {
            protocol bgp;
            community [ prepend201-3 prepend202-3 prepend203-3 prepend204-3
                prepend205-3 ];
        }
        then as-path-prepend "201 201 201";
    }
    term term4 {
        from {
            protocol bgp;

```



```

        community [ peer201 peer202 peer203 peer204 peer205 custpeer201
                    custpeer202 custpeer203 custpeer204 custpeer205 internal201
                    internal202 internal203 internal204 internal205 ];
    }
    then reject;
}
then next policy;
}
}

```

Configuring the Extended Communities Attribute

To configure extended communities, include the `community` statement:

```

community name {
    invert-match;
    members [ community-ids ];
}

```

You can include this statement at the following hierarchy levels:

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

To include the community in a routing policy, include the `community` condition in the `from` statement at the [edit policy-options policy-statement *policy-name* term *term-name*] hierarchy level:

```

[edit policy-options]
community name members [ community-ids ];
policy-statement policy-name {
    term term-name {
        from {
            community name;
        }
    }
}
}

```

name identifies one or more routers in the BGP extended community.

community-ids identifies the type of extended community in the following format:

type:administrator:assigned-number

type is the type of extended community and can be either a **bandwidth**, **target**, **origin**, **domain-id**, **src-as**, **rt-import** community, or a 16-bit number that identifies a specific BGP extended community. The **bandwidth** community sets up the bandwidth extended community. The **target** community identifies the destination to which the route is going. The **origin** community identifies where the route originated. The **domain-id** community identifies the Open Shortest Path First (OSPF) domain from which the route originated. The **src-as** community identifies the autonomous system from which the route originated. The **rt-import** community identifies the route to install in the routing table.



NOTE: For `src-as`, you can specify only an AS number and not an IP address. For `rt-import`, you can specify only an IP address and not an AS number.

administrator is the administrator. It is either an AS number or an IP version 4 (IPv4) address prefix, depending on the type of extended community.

assigned-number identifies the local provider.

Beginning with JUNOS Release 9.1, the AS number range has been extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. In plain-number format, you can configure a value in the range from 1 through 4,294,967,295. To configure an extended community for a *target* or an *origin* that includes a 4-byte AS number in the plain-number format, you must append the letter “L” to the end of number. For example, to configure a target for the 4-byte AS number 334,324 with an assigned number of 132, include the following statement: **target:334324L:132**.

Beginning with JUNOS Release 9.2, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *< 16-bit high-order value in decimal > . < 16-bit low-order value in decimal >*. For example, the 4-byte AS number of 65546 in plain-number format is represented as 1.10 in the AS-dot notation format. The *target* and *origin* extended communities support the AS-dot notation format for 4-byte AS numbers.

For information about the community attribute, see “Configuring the Community Attribute” on page 98.

Examples: Configuring the Extended Communities Attribute

Configure a target community with an administrative field of 10458 and an assigned number of 20:

```
[edit]
policy-options {
  community test-a members [ target:10458:20 ];
}
```

Configure a target community with an administrative field of 1.1.1.1 and an assigned number of 20:

```
[edit]
policy-options {
  community test-a members [ target:1.1.1.1:20 ];
}
```

Configure an origin community with an administrative field of 1.1.1.1 and an assigned number of 20:

```
[edit]
policy-options {
  community test-a members [ origin:1.1.1.1:20 ];
}
```

Inverting Community Matches

To invert the results of the community expression matching, include the `invert-match` statement:

```
invert-match;
```

You can include this statement at the following hierarchy levels:

- [edit policy-options community *name*]
- [edit logical-systems *logical-system-name* policy-options community *name*]

Configuring Link Bandwidth

To configure the link bandwidth extended communities attribute, include the `bandwidth` statement:

```
bandwidth:as:bandwidth;
```

You can include this statement at the following hierarchy levels:

- [edit policy-options community *name* members]
- [edit logical-systems *logical-system-name* policy-options community *name* members]

Specifying link bandwidth allows you to distribute traffic unequally among different BGP paths.



NOTE: The link bandwidth attribute does not work concurrently with per-prefix load balancing.

How Communities Are Evaluated

The policy framework software evaluates communities as follows:

- Each route is evaluated against each named community in a routing policy **from** statement. If a route matches one of the named communities in the **from** statement, the evaluation of the current term continues. If a route does not match, the evaluation of the current term ends.
- The route is evaluated against each member of a named community. The evaluation of all members must be successful for the named community evaluation to be successful.
- Each member in a named community is either a literal community value or a regular expression. Each member is evaluated against each community associated with the route. (Communities are an unordered property of a route. For example, 1:2 3:4 is the same as 3:4 1:2.) Only one community from the route is required to match for the member evaluation to be successful.
- Community regular expressions are evaluated on a character-by-character basis. For example, if a route contains community 1234:5678, the regular expressions

see nine discrete characters, including the colon (:), instead of two sets of numbers (1234 and 5678) separated by a colon. For example:

```
[edit]
policy-options {
  policy-statement one {
    from {
      community [comm-one comm-two];
    }
  }
  community comm-one members [ 1:2 "^4:(5|6)$" ];
  community comm-two members [ 7:8 9:10 ];
}
```



NOTE: If a community member is a regular expression, a string match is made rather than a numeric match.

- To match routing policy **one**, the route must match either **comm-one** or **comm-two**.
- To match **comm-one**, the route must have a community that matches **1:2** and a community that matches **4:5** or **4:6**.
- To match **comm-two**, the route must have a community that matches **7:8** and a community that matches **9:10**.

Configuring Prefix Lists

A *prefix list* is a named list of IP addresses. You can specify an exact match with incoming routes and apply a common action to all matching prefixes in the list.



NOTE: Because the configuration of prefix lists includes setting up prefixes and prefix lengths, we strongly recommend that you have a thorough understanding of IP addressing, including supernetting, before proceeding with the configuration.

This section includes the following information:

- Prefix List and Route List Differences on page 110
- Defining Prefix Lists on page 111
- How a Prefix List Is Evaluated on page 112
- Configuring a Prefix List Filter on page 112
- Example: Configuring a Prefix List on page 113

Prefix List and Route List Differences

A prefix list functions like a route list that contains multiple instances of the **exact** match type only. The differences between these two extended match conditions are summarized in Table 23 on page 111.

Table 23: Prefix List and Route List Differences

Feature	Prefix List	Route Lists
Action	Can specify action in a then statement only. These actions are applied to all prefixes that match the term.	Can specify action that is applied to a particular prefix in a route-filter match condition in a from statement, or to all prefixes in the list using a then statement.

Defining Prefix Lists

You can create a named prefix list and include it in a routing policy with the **prefix-list** match condition (described in Table 14 on page 42).

To define a prefix list, include the **prefix-list** statement:

```
prefix-list prefix-list-name {
  apply-path path;
  ip-addresses;
}
```

You can include this statement at the following hierarchy levels:

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

You can use the **apply-path** statement to include all prefixes pointed to by a defined path, or you can specify one or more addresses, or both.

To include a prefix list in a routing policy, specify the **prefix-list** match condition in the **from** statement at the [edit policy-options policy-statement *policy-name* term *term-name*] hierarchy level:

```
[edit policy-options policy-statement policy-name term term-name]
from {
  prefix-list prefix-list-name;
}
then actions;
```

name identifies the prefix list. It can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").

ip-addresses are the IPv4 or IP version 6 (IPv6) prefixes specified as *prefix/prefix-length*. If you omit *prefix-length* for an IPv4 prefix, the default is */32*. If you omit *prefix-length* for an IPv6 prefix, the default is */128*. Prefixes specified in a **from** statement must be either all IPv4 addresses or all IPv6 addresses.



NOTE: You cannot apply actions to individual prefixes in the list.

You can specify the same prefix list in the **from** statement of multiple routing policies or firewall filters. For information about firewall filters, see “Configuring Firewall Filters” on page 165.

Use the **apply-path** statement to configure a prefix list comprising all IP prefixes pointed to by a defined path. This eliminates most of the effort required to maintain a group prefix list.

The path consists of elements separated by spaces. Each element matches a configuration keyword or an identifier, and you can use wildcards to match more than one identifier. Wildcards must be enclosed in angle brackets, for example, `< * >`.



NOTE: When you use **apply-path** to define a prefix list, you can also use the same prefix list in a policy statement.

For examples of configuring a prefix list, see “Example: Configuring a Prefix List” on page 113; for examples of configuring a firewall filter, see “Configuring Firewall Filters” on page 165.

How a Prefix List Is Evaluated

During prefix list evaluation, the policy framework software performs a *longest-match lookup*, which means that the software searches for the prefix in the list with the longest length. The order in which you specify the prefixes, from top to bottom, does not matter. The software then compares a route’s source address to the longest prefix.

You can use prefix list qualifiers for prefixes contained in a prefix list by configuring a prefix list filter. For more information, see “Configuring Prefix Lists” on page 110.

If a match occurs, the evaluation of the current term continues. If a match does not occur, the evaluation of the current term ends.



NOTE: If you specify multiple prefixes in the prefix list, only one prefix must match for a match to occur. The prefix list matching is effectively a logical OR operation.

Configuring a Prefix List Filter

A prefix list filter allows you to apply prefix list qualifiers to a list of prefixes within a prefix list. The prefixes within the list are evaluated using the specified qualifiers. You can configure multiple prefix list filters under the same policy term.

To configure a prefix list filter, include the **prefix-list-filter** statement at the **[edit policy-options policy-statement *policy-name* from]** hierarchy level:

```
[edit policy-options policy-statement policy-name
from {
  prefix-list-filter prefix-list-name match-type actions;
```

```
}
```

The *prefix-list-name* option is the name of the prefix list to be used for evaluation. You can specify only one prefix list.

The *match-type* option is the type of match to apply to the prefixes in the prefix list. It can be one of the match types listed in Table 24 on page 113.

The *actions* option is the action to take if the prefix list matches. It can be one or more of the actions listed in Table 15 on page 48 and Table 16 on page 49.

Table 24: Route List Match Types for a Prefix List Filter

Match Type	Match Condition
exact	The route shares the same most-significant bits (described by <i>prefix-length</i>), and <i>prefix-length</i> is equal to the route's prefix length.
longer	The route shares the same most-significant bits (described by <i>prefix-length</i>), and <i>prefix-length</i> is greater than the route's prefix length.
orlonger	The route shares the same most-significant bits (described by <i>prefix-length</i>), and <i>prefix-length</i> is equal to or greater than the route's prefix length.

Example: Configuring a Prefix List

The following example accepts and rejects traffic from sites specified using prefix lists:

```
[edit]
policy-options {
  policy-statement prefix-list-policy {
    term ok-sites {
      from {
        prefix-list known-ok-sites;
      }
      then accept;
    }
    term reject-bcasts {
      from {
        prefix-list known-dir-bcast-sites;
      }
      then reject;
    }
  }
}
[edit]
policy-options {
  prefix-list known-ok-sites {
    172.16.0.3;
    10.10.0.0/16;
    192.168.12.0/24;
  }
}
```

```
[edit]
policy-options {
  prefix-list known-dir-bcast-sites {
    10.3.4.6;
    10.2.0.0/16;
    192.168.1.0/24;
  }
}
```

Configuring Route Lists

A *route list* is a collection of destination prefixes. When specifying a prefix, you can specify an exact match with a particular route or a less precise match. You can configure either a common action that applies to the entire list or an action associated with each prefix.



NOTE: Because the configuration of route lists includes setting up prefixes and prefix lengths, we strongly recommend that you have a thorough understanding of IP addressing, including supernetting, before proceeding with the configuration.

It is also important to understand how a route list is evaluated, particularly if the route list includes multiple **route-filter** options in a **from** statement. We strongly recommend that you read “How a Route List Is Evaluated” on page 117 before proceeding with the configuration. Not fully understanding the evaluation process could result in faulty configuration and unexpected results.

This section discusses the following topics:

- Defining Route Lists on page 114
- How a Route List Is Evaluated on page 117
- Examples: Configuring Route Lists on page 119

Defining Route Lists

To specify a route list, include one or more **route-filter** or **source-address-filter** options in the **from** statement of the **policy-statement** at the [edit policy-options policy-statement *policy-name* term *term-name*] hierarchy level:

```
[edit policy-options policy-statement policy-name term term-name]
from {
  route-filter prefix match-type <actions>;
  source-address-filter source-prefix match-type <actions>;
}
then actions;
```

The **route-filter** option is typically used to match prefixes of any type except for unicast source addresses.

The **source-address-filter** option is typically used to match unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments.

source-prefix is the IPv4 or IPv6 prefix specified as *prefix/prefix-length*. If you omit *prefix-length* for an IPv4 prefix, the default is /32. If you omit *prefix-length* for an IPv6 prefix, the default is /128. Prefixes specified in a **from** statement must be either all IPv4 addresses or all IPv6 addresses.

match-type is the type of match to apply to the destination prefix. It can be one of the match types listed in Table 25 on page 115. For examples of the match types and the results when presented with various routes, see Table 26 on page 116.

actions is the action to take if the destination prefix matches. It can be one or more of the actions listed in Table 15 on page 48 and Table 16 on page 49.

In route lists, you can specify actions in two ways:

- In the **route-filter** or **source-address-filter** option—These actions are taken immediately after a match occurs, and the **then** statement is not evaluated.
- In the **then** statement—These actions are taken after a match occurs and if an action is not specified in the **route-filter** or **source-address-filter** option.

The **upto** and **prefix-length-range** match types are similar in that both specify the most-significant bits and provide a range of prefix lengths that can match. The difference is that **upto** allows you to specify an upper limit only for the prefix length range, whereas **prefix-length-range** allows you to specify both lower and upper limits.

For more examples of these route list match types, see “Examples: Configuring Route Lists” on page 119.

Table 25: Route List Match Types for a Prefix List

Match Type	Match Condition
exact	The route shares the same most-significant bits (described by <i>prefix-length</i>), and <i>prefix-length</i> is equal to the route's prefix length.
longer	The route shares the same most-significant bits (described by <i>prefix-length</i>), and <i>prefix-length</i> is greater than the route's prefix length.
orlonger	The route shares the same most-significant bits (described by <i>prefix-length</i>), and <i>prefix-length</i> is equal to or greater than the route's prefix length.
prefix-length-range <i>prefix-length2–prefix-length3</i>	The route shares the same most-significant bits (described by <i>prefix-length</i>), and the route's prefix length falls between <i>prefix-length2</i> and <i>prefix-length3</i> , inclusive.

Table 25: Route List Match Types for a Prefix List (*continued*)

Match Type	Match Condition
through <i>destination-prefix</i>	<p>All the following are true:</p> <ul style="list-style-type: none"> ■ The route shares the same most-significant bits (described by <i>prefix-length</i>) of the first destination prefix. ■ The route shares the same most-significant bits (described by <i>prefix-length</i>) of the second destination prefix for the number of bits in the prefix length. ■ The number of bits in the route's prefix length is less than or equal to the number of bits in the second prefix. <p>You do not use the through match type in most routing policy configurations. (For an example, see “Examples: Configuring Route Lists” on page 119.)</p>
upto <i>prefix-length2</i>	The route shares the same most-significant bits (described by <i>prefix-length</i>) and the route's prefix length falls between <i>prefix-length</i> and <i>prefix-length2</i> .

Table 26: Match Type Examples

Prefix	192.168/16 exact	192.168/16 longer	192.168/16 orlonger	192.168/16 upto /24	192.168/16 through 192.168.16/20	192.168/16 prefix length range /18–/20
10.0.0.0/8	–	–	–	–	–	–
192.168.0.0/16	Match	–	Match	Match	Match	–
192.168.0.0/17	–	Match	Match	Match	Match	–
192.168.0.0/18	–	Match	Match	Match	Match	Match
192.168.0.0/19	–	Match	Match	Match	Match	Match
192.168.4.0/24	–	Match	Match	Match	–	–
192.168.5.4/30	–	Match	Match	–	–	–
192.168.12.4/30	–	Match	Match	–	–	–
192.168.12.128/32	–	Match	Match	–	–	–
192.168.16.0/20	–	Match	Match	Match	Match	Match
192.168.192.0/18	–	Match	Match	Match	–	Match
192.168.224.0/19	–	Match	Match	Match	–	Match
10.169.1.0/24	–	–	–	–	–	–
10.170.0.0/16	–	–	–	–	–	–

How a Route List Is Evaluated

During route list evaluation, the policy framework software compares each route's source address with the destination prefixes in the route list. The evaluation occurs in two steps:

1. The policy framework software performs a *longest-match lookup*, which means that the software searches for the prefix in the list with the longest length.

The longest-match lookup considers the prefix and prefix length only and not the match type. The following sample route list illustrates this point:

```
from {
  route-filter 192.168.0.0/14 upto /24 reject;
  route-filter 192.168.0.0/15 exact;
}
then accept;
```

The longest match is the second route-filter, **192.168.0.0/15**, which is based on prefix and prefix length only.

2. Once an incoming route matches a prefix (longest first), the following occur:
 - The route filter stops evaluating other prefixes, even if the match type fails.
 - The software examines the match type and action associated with that prefix.

In Step 1, if route **192.168.1.0/24** were evaluated, it would fail to match. It matches the longest prefix of **192.168.0.0/15**, but it does not match **exact**. The route filter is finished because it matched a prefix, but the result is a failed match because the match type failed.

If a match occurs, the action specified with the prefix is taken. If an action is not specified with the prefix, the action in the **then** statement is taken. If neither action is specified, the software evaluates the next term or routing policy, if present, or takes the **accept** or **reject** action specified by the default policy. For more information about the default routing policies, see “Default Routing Policies and Actions” on page 19.



NOTE: If you specify multiple prefixes in the route list, only one prefix needs to match for a match to occur. The route list matching is effectively a logical OR operation.

If a match does not occur, the software evaluates the next term or routing policy, if present, or takes the **accept** or **reject** action specified by the default policy.

For example, compare the prefix **192.168.254.0/24** against the following route list:

```
route-filter 192.168.0.0/16 orlonger;
route-filter 192.168.254.0/23 exact;
```

The prefix **192.168.254.0/23** is determined to be the longest prefix. When the software evaluates **192.168.254.0/24** against the longest prefix, a match occurs

(192.168.254.0/24 is a subset of 192.168.254.0/23). Because of the match between 192.168.254.0/24 and the longest prefix, the evaluation continues. However, when the software evaluates the match type, a match does not occur between 192.168.254.0/24 and 192.168.254.0/23 exact. The software concludes that the term does not match and goes on to the next term or routing policy, if present, or takes the **accept** or **reject** action specified by the default policy.

The following sections describe the following topics:

- How Prefix Order Affects Route List Evaluation on page 118
- Common Configuration Problem with the Longest-Match Lookup on page 118

How Prefix Order Affects Route List Evaluation

The order in which the prefixes are specified (from top to bottom) typically does not matter, because the policy framework software scans the route list looking for the longest prefix during evaluation. An exception to this rule is when you use the same destination prefix multiple times in a list. In this case, the order of the prefixes is important, because the list of identical prefixes is scanned from top to bottom, and the first match type that matches the route applies.

In the following example, different match types are specified for the same prefix. The route 0.0.0.0/0 would be rejected, the route 0.0.0.0/8 would be marked with **next-hop self**, and the route 0.0.0.0/25 would be rejected.

```
route-filter 0.0.0.0/0 upto /7 reject;
route-filter 0.0.0.0/0 upto /24 next-hop self;
route-filter 0.0.0.0/0 orlonger reject;
```

Common Configuration Problem with the Longest-Match Lookup

A common problem when defining a route list is including a shorter prefix that you want to match with a longer, similar prefix in the same list. For example, imagine that the prefix 192.168.254.0/24 is compared against the following route list:

```
route-filter 192.168.0.0/16 orlonger;
route-filter 192.168.254.0/23 exact;
```

Because the policy framework software performs longest-match lookup, the prefix 192.168.254.0/23 is determined to be the longest prefix. An exact match does not occur between 192.168.254.0/24 and 192.168.254.0/23 exact. The software determines that the term does not match and goes on to the next term or routing policy, if present, or takes the **accept** or **reject** action specified by the default policy. (For more information about the default routing policies, see “Default Routing Policies and Actions” on page 19.) The shorter prefix 192.168.0.0/16 orlonger that you wanted to match is inadvertently ignored.

One solution to this problem is to remove the prefix 192.168.0.0/16 orlonger from the route list in this term and move it to a previous term where it is the only prefix or the longest prefix in the list.

Examples: Configuring Route Lists

The examples in this section show only fragments of routing policies. Normally, you would combine these fragments with other terms or routing policies.

In all examples, remember that the following actions apply to nonmatching routes:

- Evaluate next term, if present.
- Evaluate next policy, if present.
- Take the accept or reject action specified by the default policy. For more information about the default routing policies, see “Default Routing Policies and Actions” on page 19.

The following examples show how to configure route lists for various purposes:

- Example: Rejecting Routes with Specific Destination Prefixes and Mask Lengths on page 119
- Example: Rejecting Routes with a Mask Length Greater than Eight on page 119
- Example: Rejecting Routes with Mask Length Between 26 and 29 on page 120
- Example: Rejecting Routes from Specific Hosts on page 120
- Example: Accepting Routes with a Defined Set of Prefixes on page 121
- Example: Rejecting Routes with a Defined Set of Prefixes on page 121
- Example: Rejecting Routes with Prefixes Longer than 24 Bits on page 122
- Example: Rejecting PIM Multicast Traffic Joins on page 122
- Example: Rejecting PIM Traffic on page 122

Example: Rejecting Routes with Specific Destination Prefixes and Mask Lengths

Reject routes with a destination prefix of 0.0.0.0 and a mask length from 0 through 8, and accept all other routes:

```
[edit]
policy-options {
  policy-statement policy-statement from-hall2 {
    term 1 {
      from {
        route-filter 0.0.0.0/0 upto /8 reject;
      }
    }
    then accept;
  }
}
```

Example: Rejecting Routes with a Mask Length Greater than Eight

Reject routes with a mask of /8 and greater (that is, /8, /9, /10, and so on) that have the first 8 bits set to 0 and accept routes less than 8 bits in length:

```
[edit]
policy-options {
  policy-statement from-hall3 {
    term term1 {
      from {
        route-filter 0/0 upto /7 accept;
        route-filter 0/8 orlonger;
      }
      then reject;
    }
  }
}
```

Example: Rejecting Routes with Mask Length Between 26 and 29

Reject routes with the destination prefix of 192.168.10/24 and a mask between /26 and /29 and accept all other routes:

```
[edit]
policy-options {
  policy-statement from-customer-a {
    term term1 {
      from {
        route-filter 192.168.10/24 prefix-length-range /26-/29 reject;
        route-filter 0/0;
      }
      then accept;
    }
  }
}
```

Example: Rejecting Routes from Specific Hosts

Reject a range of routes from specific hosts, and accept all other routes:

```
[edit]
policy-options {
  policy-statement hosts-only {
    from {
      route-filter 10.125.0.0/16 upto /31 reject;
      route-filter 0/0;
    }
    then accept;
  }
}
```

You do not use the **through** match type in most routing policy configurations. You should think of **through** as a tool to group a contiguous set of exact matches. For example, instead of specifying four exact matches:

```
from route-filter 0.0.0.0/1 exact
from route-filter 0.0.0.0/2 exact
from route-filter 0.0.0.0/3 exact
from route-filter 0.0.0.0/4 exact
```

You could represent them with the following single match:

```
from route-filter 0.0.0.0/1 through 0.0.0.0/4
```

Example: Accepting Routes with a Defined Set of Prefixes

Explicitly accept a limited set of prefixes (in the first term) and reject all others (in the second term):

```
policy-options {
  policy-statement internet-in {
    term 1 {
      from {
        route-filter 192.168.231.0/24 exact accept;
        route-filter 192.168.244.0/24 exact accept;
        route-filter 192.168.198.0/24 exact accept;
        route-filter 192.168.160.0/24 exact accept;
        route-filter 192.168.59.0/24 exact accept;
      }
    }
    term 2 {
      then {
        reject;
      }
    }
  }
}
```

Example: Rejecting Routes with a Defined Set of Prefixes

Reject a few groups of prefixes, and accept the remaining prefixes:

```
[edit policy-options]
policy-statement drop-routes {
  term 1{
    from { # first, reject a number of prefixes:
      route-filter default exact reject; # reject 0.0.0.0/0 exact
      route-filter 0.0.0.0/8 orlonger reject; # reject prefix 0, mask /8 or longer
      route-filter 10.0.0.0/8 orlonger reject; # reject loopback addresses
    }
    route-filter 10.105.0.0/16 exact { # accept 10.105.0.0/16
      as-path-prepend "1 2 3";
      accept;
    }
    route-filter 192.0.2.0/24 orlonger reject; # reject test network packets
    route-filter 224.0.0.0/3 orlonger reject; # reject multicast and higher
    route-filter 0.0.0.0/0 upto /24 accept; # accept everything up to /24
    route-filter 0.0.0.0/0 orlonger accept; # accept everything else
  }
}
```

Example: Rejecting Routes with Prefixes Longer than 24 Bits

Reject all prefixes longer than 24 bits. You would install this routing policy in a sequence of routing policies in an **export** statement. The first term in this filter passes on all routes with a prefix length of up to 24 bits. The second, unnamed term rejects everything else.

```
[edit policy-options]
policy-statement 24bit-filter {
  term acl20 {
    from {
      route-filter 0.0.0.0/0 upto /24;
    }
    then next policy;
  }
  then reject;
}
```

If, in this example, you were to specify **route-filter 0.0.0.0/0 upto /24 accept**, matching prefixes would be accepted immediately and the next routing policy in the **export** statement would never get evaluated.

If you were to include the **then reject** statement in the term **acl20**, prefixes greater than 24 bits would never get rejected because the policy framework software, when evaluating the term, would move on to evaluating the next statement before reaching the **then reject** statement.

Example: Rejecting PIM Multicast Traffic Joins

Configure a routing policy for rejecting Protocol Independent Multicast (PIM) multicast traffic joins for a source destination prefix from a neighbor:

```
[edit]
policy-options {
  policy-statement join-filter {
    from {
      neighbor 10.14.12.20;
      source-address-filter 10.83.0.0/16 orlonger;
    }
    then reject;
  }
}
```

Example: Rejecting PIM Traffic

Configure a routing policy for rejecting PIM traffic for a source destination prefix from an interface:

```
[edit]
policy-options {
  policy-statement join-filter {
    from {
      interface so-1/0/0.0;
      source-address-filter 10.83.0.0/16 orlonger;
    }
  }
}
```



```

    }
    then reject;
  }
}

```

The following routing policy qualifiers apply to PIM:

- **interface**—Interface over which a join is received
- **neighbor**—Source from which a join originates
- **route-filter**—Group address
- **source-address-filter**—Source address for which to reject a join

For more information about importing a PIM join filter in a PIM protocol definition, see the *JUNOS Multicast Protocols Configuration Guide*.

Configuring Subroutines

You can use a routing policy called from another routing policy as a match condition. This process makes the called policy a *subroutine*.

The following sections discuss configuring subroutines and provide the following example:

- Defining Subroutines on page 123
- Example: Configuring a Subroutine on page 126

Defining Subroutines

To configure a subroutine in a routing policy to be called from another routing policy, create the subroutine and specify its name using the **policy** match condition (described in Table 14 on page 42) in the **from** or **to** statement of another routing policy:

```

[edit]
policy-options {
  policy-statement subroutine-policy-name {
    term term-name {
      from {
        match-conditions;
        route-filter destination-prefix match-type <actions>;
        source-address-filter destination-prefix match-type <actions>;
        prefix-list name;
      }
      to {
        match-conditions;
      }
      then actions;
    }
  }
}
policy-options {
  policy-statement policy-name {
    term term-name {

```

```

        from {
            policy subroutine-policy-name;
        }
        to {
            policy subroutine-policy-name;
        }
        then actions;
    }
}

```



NOTE: Do not evaluate a routing policy within itself. If you attempt to do so, no prefixes will ever match the routing policy.

The action specified in a subroutine is used to provide a match condition to the calling policy. If the subroutine specifies an action of accept, the calling policy considers the route to be a match. If the subroutine specifies an action of reject, the calling policy considers the route not to match. If the subroutine specifies an action that is meant to manipulate the route characteristics, the changes are made. For more details about the subroutine evaluation, see “How a Routing Policy Subroutine Is Evaluated” on page 30.

Termination Actions

A subroutine with particular statements can behave differently from a routing policy that contains the same statements. With a subroutine, you must remember that the possible termination actions of accept or reject specified by the subroutine or the default policy can greatly affect the expected results. (For more information about default routing policies, see “Default Routing Policies and Actions” on page 19.)

In particular, you must consider what happens if a match does not occur with routes specified in a subroutine and if the default policy action that is taken is the action that you expect and want.

For example, imagine that you are a network administrator at an Internet service provider (ISP) that provides service to Customer A. You have configured several routing policies for the different classes of neighbors that Customer A presents on various links. To save time maintaining the routing policies for Customer A, you have configured a subroutine that identifies their routes and various routing policies that call the subroutine, as shown below:

```

[edit]
policy-options {
  policy-statement customer-a-subroutine {
    from {
      route-filter 10.1/16 exact;
      route-filter 10.5/16 exact;
      route-filter 192.168.10/24 exact;
    }
    then accept;
  }
}

```

```

policy-options {
  policy-statement send-customer-a-default {
    from {
      policy customer-a-subroutine;
    }
    then {
      set metric 500;
      accept;
    }
  }
}
policy-options {
  policy-statement send-customer-a-primary {
    from {
      policy customer-a-subroutine;
    }
    then {
      set metric 100;
      accept;
    }
  }
}
policy-options {
  policy-statement send-customer-a-secondary {
    from {
      policy customer-a-subroutine;
    }
    then {
      set metric 200;
      accept;
    }
  }
}
protocols {
  bgp {
    group customer-a {
      export send-customer-a-default;
      neighbor 10.1.1.1;
      neighbor 10.1.2.1;
      neighbor 10.1.3.1 {
        export send-customer-a-primary;
      }
      neighbor 10.1.4.1 {
        export send-customer-a-secondary;
      }
    }
  }
}

```

The following results occur with this configuration:

- The group-level **export** statement resets the metric to 500 when advertising all BGP routes to neighbors 10.1.1.1 and 10.1.2.1 rather than just the routes that match the subroutine route filters.

- The neighbor-level **export** statements reset the metric to **100** and **200** when advertising all BGP routes to neighbors **10.1.3.1** and **10.1.4.1**, respectively, rather than just the BGP routes that match the subroutine route filters.

These unexpected results occur because the subroutine policy does not specify a termination action for routes that do not match the route filter and therefore, the default BGP export policy of accepting all BGP routes is taken.

If the statements included in this particular subroutine had been contained within the calling policies themselves, only the desired routes would have their metrics reset.

This example illustrates the differences between routing policies and subroutines and the importance of the termination action in a subroutine. Here, the default BGP export policy action for the subroutine was not carefully considered. A solution to this particular example is to add one more term to the subroutine that rejects all other routes that do not match the route filters:

```
[edit]
policy-options {
  policy-statement customer-a-subroutine {
    term accept-exact {
      from {
        route-filter 10.1/16 exact;
        route-filter 10.5/16 exact;
        route-filter 192.168.10/24 exact;
      }
      then accept;
    }
    term reject-others {
      then reject;
    }
  }
}
```

Termination action strategies for subroutines in general include the following:

- Depend upon the default policy action to handle all other routes.
- Add a term that accepts all other routes. (Also see “Side Effects of Omitting the “from” Statement from an Export Policy” on page 64.)
- Add a term that rejects all other routes.

The option that you choose depends upon what you want to achieve with your subroutine. Plan your subroutines carefully.

Example: Configuring a Subroutine

Create the subroutine **is-customer** and call it from the routing policies **export-customer** and **import-customer**. In **import-customer**, the action is taken only on routes that match the route filters defined in **is-customer**.

```
[edit]
policy-options {
```

```

policy-statement is-customer {
  term match-customer {
    from {
      route-filter 10.100.1.0/24 exact;
      route-filter 10.186.100.0/24 exact;
    }
    then accept;
  }
  term drop-others {
    then reject;
  }
}
policy-statement export-customer {
  from policy is-customer;
  then accept;
}
policy-statement import-customer {
  from {
    protocol bgp;
    policy is-customer;
  }
  then {
    local-preference 10;
    accept;
  }
}
}

```

Configuring the Condition Statement

In addition to the well-established use of match conditions in the **from** statement of a policy, you can also configure the **condition** statement to define conditions that will be used in the policy evaluation:

```

condition condition-name {
  if-route-exists address table table-name;
}

```

You can include this statement at the following hierarchy levels:

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

To define a policy condition based on the existence of routes in specific tables for use in BGP export policies, specify a name for the condition and include the following options:

- **if-route-exists** *address*—Specify the address of the route in question.
- **table** *table-name*—Specify a routing table.

You can then add the defined condition to the **from** statement of a policy:

```

policy-options {

```

```
policy-statement policy-name {  
  term 1 {  
    from {  
      protocols bgp;  
      condition condition-name;  
    }  
    then {  
      accept;  
    }  
  }  
  ...  
}
```

The **condition** statement is available on all platforms, but is limited to use in BGP export policies. To view the configured policy conditions and their associated routing tables and dependent routes, issue the **show policy conditions** operational mode command; for more information, see the *JUNOS Routing Protocols and Policies Command Reference*.

Chapter 6

Extended Actions Configuration

This chapter describes how to configure extended actions for a routing policy. These extended actions include criteria that manipulate the route characteristics and are more complex than standard actions. You should have an in-depth understanding of extended actions before configuring them. For a complete list of the actions that manipulate route characteristics, see Table 16 on page 49.

This chapter provides information about the following routing policy actions configuration tasks:

- Configuring the AS Path Prepend Action on page 129
- Configuring the AS Path Expand Action on page 130
- Configuring the Class-of-Service Action on page 130
- Configuring the Damping Action on page 130
- Configuring the Load-Balance Per-Packet Action on page 136

Configuring the AS Path Prepend Action

You can *prepend* or add one or more autonomous system (AS) numbers at the beginning of an AS path. The AS numbers are added after the local AS number has been added to the path. Prepending an AS path makes a shorter AS path look longer and therefore less preferable to the Border Gateway Protocol (BGP).

Beginning with JUNOS Release 9.1, the range that you can configure for AS numbers has been extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for 4-octet AS Number Space*. The range is now from 1 through 4,294,967,295. The JUNOS software continues to support 2-byte AS numbers. For more information about configuring the AS number for the router and BGP, see the *JUNOS Routing Protocols Configuration Guide*.

For example, from AS 1, there are two equal paths (through AS 2 and AS 3) to reach AS 4. You might want packets from certain sources to use the path through AS 2. Therefore, you must make the path through AS 3 look less preferable so that BGP chooses the path through AS 2. In AS 1, you can prepend multiple AS numbers:

```
[edit]
policy-options {
  policy-statement as-path-prepend {
    term prepend {
      from {
        route-filter 192.168.0.0/16 orlonger;
```

```

        route-filter 172.16.0.0/12 orlonger;
        route-filter 10.0.0.0/8 orlonger;
    }
    then as-path-prepend "1 1 1 1";
}
}
}

```

Configuring the AS Path Expand Action

You can expand or add one or more AS numbers to an AS sequence. The AS numbers are added before the local AS number has been added to the path. Expanding an AS path makes a shorter AS path look longer and therefore less preferable to BGP. The last AS number in the existing path is extracted and prepended *n* times, where *n* is a number from 1 through 32. This is similar to the AS path prepend action, except that the AS path expand action adds an arbitrary sequence of AS numbers.

For example, from AS 1, there are two equal paths (through AS 2 and AS 3) to reach AS 4. You might want packets from certain sources to use the path through AS 2. Therefore, you must make the path through AS 3 look less preferable so that BGP chooses the path through AS 2. In AS 1, you can expand multiple AS numbers.

```

[edit]
policy-options {
  policy-statement as-path-expand {
    term expand {
      from {
        route-filter 192.168.0.0/16 orlonger;
        route-filter 172.16.0.0/12 orlonger;
        route-filter 10.0.0.0/8 orlonger;
      }
      then as-path-expand last-as count 4;
    }
  }
}

```

For routes from AS 2, this makes the route look like **1 2 2 2 2 2** when advertised, where **1** is from AS 1, the **2** from AS 2 is prepended four times, and the final **2** is the original 2 received from the neighbor router.

Configuring the Class-of-Service Action

For information about class of service (CoS), see the *JUNOS Class of Service Configuration Guide*.

Configuring the Damping Action

BGP *route flapping* describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability information. BGP *flap damping* is a way to reduce the number of update messages sent between BGP peers, thereby reducing the load on these peers without adversely affecting the route convergence time.

Flap damping reduces the number of update messages by marking routes as ineligible for selection as the active or preferable route. Doing this leads to some delay, or *suppression*, in the propagation of route information, but the result is increased network stability. You typically apply flap damping to external BGP (EBGP) routes (that is, to routes in different ASs). You can also apply it within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.)

BGP flap damping is defined in RFC 2439, *BGP Route Flap Damping*.

To effect changes to the default BGP flap damping values, you define actions by creating a named set of damping parameters and including it in a routing policy with the **damping** action (described in Table 16 on page 49). For the damping routing policy to work, you also must enable BGP route flap damping.

The following sections discuss the following topics:

- Configuring Flap Damping Parameters on page 131
- Defining Damping Action on page 133
- Enabling BGP Route Flap Damping on page 134
- Disabling Damping by Prefix on page 134
- Example: Configuring BGP Flap Damping on page 134

Configuring Flap Damping Parameters

To define damping parameters, include the **damping** statement:

```
damping name {
  disable;
  half-life minutes;
  max-suppress minutes;
  reuse number;
  suppress number;
}
```

You can include this statement at the following hierarchy levels:

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

The name identifies the group of damping parameters. It can contain letters, numbers, and hyphens (-) and can be up to 255 characters. To include spaces in the name, enclose the entire name in quotation marks (" ").

You can specify one or more of the damping parameters described in Table 27 on page 132.

Table 27: Damping Parameters

Damping Parameter	Description	Default	Possible Values
half-life <i>minutes</i>	Decay half-life, in minutes	15 minutes	1 through 45 minutes
max-suppress <i>minutes</i>	Maximum hold-down time, in minutes	60 minutes	1 through 720 minutes
reuse	Reuse threshold	750 (unitless)	1 through 20,000 (unitless)
suppress	Cutoff (suppression) threshold	3000 (unitless)	1 through 20,000 (unitless)

If you do not specify one or more of the damping parameters, the default value of the parameter is used.

To understand how to configure these parameters, you need to understand how damping suppresses routes. How long a route can be suppressed is based on a *figure of merit*, which is a value that correlates to the probability of future instability of a route. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time.

A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes. With each incident of instability, the value increases as follows:

- Route is withdrawn—1000
- Route is readvertised—1000
- Route's path attributes change—500



NOTE: Other vendors' implementations for figure-of-merit increase the value only when a route is withdrawn. The JUNOS implementation for figure-of-merit increases the value for both route withdrawal and route readvertisement. To accommodate other implementations for figure-of-merit, multiply the **reuse** and **suppress** threshold values by 2.

When a route's figure-of-merit value reaches a particular level, called the *cutoff* or *suppression threshold*, the route is suppressed. If a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols. By default, a route is suppressed when its figure-of-merit value reaches 3000. To modify this default, include the **suppress** option at the [edit policy-options damping *name*] hierarchy level.

If a route has flapped, but then becomes stable so that none of the incidents listed previously occur within a configurable amount of time, the figure-of-merit value for the route decays exponentially. The default half-life is 15 minutes. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes. To

modify the default half-life, include the **half-life** option at the [edit policy-options damping *name*] hierarchy level.

A suppressed route becomes reusable when its figure-of-merit value decays to a value below a *reuse threshold*, thus allowing routes that experience transient instability to once again be considered valid. The default reuse threshold is 750. When the figure-of-merit value passes below the reuse threshold, the route once again is considered usable and can be installed in the forwarding table and exported from the routing table. To modify the default reuse threshold, include the **reuse** option at the [edit policy-options damping *name*] hierarchy level.

The maximum suppression time provides an upper bound on the time that a route can remain suppressed. The default maximum suppression time is 60 minutes. To modify the default, include the **max-suppress** option at the [edit policy-options damping *name*] hierarchy level.

A route's figure-of-merit value stops increasing when it reaches a maximum suppression threshold, which is determined based on the route's suppression threshold level, half-life, reuse threshold, and maximum hold-down time.

The merit ceiling, ϵ_c , which is the maximum merit that a flapping route can collect, is calculated using the following formula:

$$\epsilon_c \leq \epsilon_r e^{(t/\lambda) (\ln 2)}$$

ϵ_r is the figure-of-merit reuse threshold, t is the maximum hold-down time in minutes, and λ is the half-life in minutes. For example, if you use the default figure-of-merit values in this formula, but use a half-life of 30 minutes, the calculation is as follows:

$$\epsilon_c \leq 750 e^{(60/30) (\ln 2)}$$

$$\epsilon_c \leq 3000$$



NOTE: The cutoff threshold, which you configure using the **suppress** option, must be less than or equal to the merit ceiling, ϵ_c . If the configured cutoff threshold or the default cutoff threshold is greater than the merit ceiling, the route is never suppressed and damping never occurs.

To display figure-of-merit information, use the **show policy damping** command.

A route that has been assigned a figure of merit is considered to have a damping state. To display the current damping information on the router, use the **show route detail** command.

Defining Damping Action

To define the damping action, include the **damping** action and the name of the configured damping parameters either in a **then** statement or in a **route-filter** option in a **from** statement.

Enabling BGP Route Flap Damping

For information about enabling BGP route flap damping, see the *JUNOS Routing Protocols Configuration Guide*.

Disabling Damping by Prefix

Normally, you enable or disable damping on a per-peer basis. However, you can disable damping for a specific prefix received from a peer by including the **disable** option:

```
disable;
```

You can include this statement at the following hierarchy levels:

- [edit policy-options damping *name*]
- [edit logical-systems *logical-system-name* policy-options damping *name*]

Example: Disabling by Prefix

In this routing policy example, although damping is enabled for the peer, the **damping none** statement specifies that damping be disabled for prefix 3.0.0.0/8 in Policy-A. This route is not damped because the routing policy statement named **Policy-A** filters on the prefix 3.0.0.0/8 and the action points to the **damping** statement named **none**. The remaining prefixes are damped using the default parameters.

```
[edit]
policy-options {
  policy-statement Policy-A {
    from {
      route-filter 10.0.0.0/8 exact;
    }
    then damping none;
  }
  damping none {
    disable;
  }
}
```

Example: Configuring BGP Flap Damping

Enable BGP flap damping and configure damping parameters:

```
[edit]
routing-options {
  autonomous-system 666;
}
protocols {
  bgp {
    damping;
    group group1 {
      traceoptions {
        file bgp-log size 1m files 10;
      }
    }
  }
}
```

```

        flag damping;
    }
    import damp;
    type external;
    peer-as 10458;
    neighbor 192.168.2.30;
}
}
}
policy-options {
    policy-statement damp {
        from {
            route-filter 192.168.0.0/32 exact {
                damping high;
                accept;
            }
            route-filter 172.16.0.0/32 exact {
                damping medium;
                accept;
            }
            route-filter 10.0.0.0/8 exact {
                damping none;
                accept;
            }
        }
    }
}
damping high {
    half-life 30;
    suppress 3000;
    reuse 750;
    max-suppress 60;
}
damping medium {
    half-life 15;
    suppress 3000;
    reuse 750;
    max-suppress 45;
}
damping none {
    disable;
}
}

```

To display damping parameters for this configuration, use the `show policy damping` command:

```

user@host> show policy damping
Damping information for "high":
  Halflife: 30 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 3008
    Maximum decay: 24933
Damping information for "medium":
  Halflife: 15 minutes

```

```

Reuse merit: 750 Suppress/cutoff merit: 3000
Maximum suppress time: 45 minutes
Computed values:
  Merit ceiling: 6024
  Maximum decay: 12449
Damping information for "none":
Damping disabled

```

Configuring the Load-Balance Per-Packet Action

By default, when there are multiple equal-cost paths to the same destination for the active route, the JUNOS software uses a hash algorithm to choose one of the next-hop addresses to install in the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is rechosen using the hash algorithm.

You can configure the JUNOS software so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This feature is called per-packet load balancing. You can use load balancing to spread traffic across multiple paths between routers. The behavior of the load-balance per-packet function depends on the version of the Internet Processor application-specific integrated circuit (ASIC) in your routing platform.

On routing platforms with an Internet Processor ASIC, when per-packet load balancing is configured, traffic between routers with multiple paths is spread using the hash algorithm across the available interfaces. The forwarding table balances the traffic headed to a destination, transmitting it in round-robin fashion among the multiple next hops (up to a maximum of eight equal-cost load-balanced paths). The traffic is load-balanced on a per-packet basis.

On routing platforms with the Internet Processor II ASIC, when per-packet load balancing is configured, traffic between routers with multiple paths is divided into individual traffic flows (up to a maximum of 16 equal-cost load-balanced paths). Packets for each individual flow are kept on a single interface.



NOTE: You can configure per-packet load balancing to optimize VPLS traffic flows across multiple paths.

To configure the load-balance per-packet action, include the **load-balance per-packet** action in a **then** statement or a **route-filter** option in a **from** statement in a routing policy.

You must apply the routing policy to routes exported from the routing table to the forwarding table to complete the configuration. To do this, include the **export** statement:

```
export [ policy-names ];
```

You can include this statement at the following hierarchy levels:

- [edit routing-options forwarding-table]
- [edit routing-instances *routing-instance-name* routing-options forwarding-table]
- [edit logical-systems *logical-system-name* routing-options forwarding-table]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options forwarding-table]

By default, the software ignores port data when determining flows. To enable per-flow load balancing, you must set the **load-balance per-packet** action in the routing policy configuration; for more information about this action, see “Routing Policy Configuration” on page 39.

To include port data in the flow determination, include the **family inet** statement at the [edit forwarding-options hash-key] hierarchy level:

```
[edit forwarding-options hash-key]
family inet {
  layer-3;
  layer-4;
}
```

You must include the **layer-3** statement. If you omit the **layer-3** statement, the management process removes the **hash-key** statement from the configuration and the router behaves as if you specified **layer-3**.

If you include both the **layer-3** and **layer-4** statements, the router uses the following Layer 3 and Layer 4 information to load-balance:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Incoming interface index

The router recognizes packets in which all of these **layer-3** and **layer-4** parameters are identical, and ensures that these packets are sent out through the same interface. This prevents problems that might otherwise occur with packets arriving at their destination out of their original sequence.

This is appropriate behavior for Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) packets. For Internet Control Message Protocol (ICMP) packets, the field location offset is the checksum field, which makes each ping packet a separate “flow.” There are other protocols that can be encapsulated in IP that may have a varying value in the 32-bit offset. This may also be problematic because these protocols are seen as a separate flow.

By default, or if you specify only the **layer-3** statement, the router uses the incoming interface index as well as the following Layer 3 information in the packet header to load balance traffic:

- Source IP address
- Destination IP address
- Protocol

By default, Internet Protocol version 6 (IPv6) packets are automatically load-balanced based on the following Layer 3 and Layer 4 information:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Incoming interface index
- Traffic class

For more information about load balancing, see the following sections:

- Load Balancing Based on the MPLS Label Information on page 138
- Load Balancing Based on Layer 2 MAC Information on page 140
- Examples: Configuring Per-Packet Load Balancing on page 141

Load Balancing Based on the MPLS Label Information

To load-balance based on the Multiprotocol Label Switching (MPLS) label information, include the **family mpls** statement at the **[edit forwarding-options hash-key]** hierarchy level:

```
[edit forwarding-options hash-key]
family mpls {
  label-1;
  label-2;
  label-3;
  no-labels;
  no-label-1-exp;
  payload {
    ip;
  }
}
```

This is for aggregated Ethernet and aggregated SONET/SDH interfaces, as well as for multiple equal-cost MPLS next hops.

To include the first label in the hash key, include the **label-1** option. This is used for a one-label packet.

To include the first and second label in the hash key, include both the **label-1** and **label-2** options. This is used for a two-label packet. The router provides hashing on the first and second labels by default. If both labels are specified, the entire first label and the first 16 bits of the second label are hashed.

To include the third MPLS label in the hash function, include the **label-3** option at the **[edit forwarding-options hash-key family mpls]** hierarchy level. To include no MPLS labels in the hash function, include the **no-labels** option at the **[edit forwarding-options hash-key family mpls]** hierarchy level.

Hashing may include IP addresses to provide better distribution of traffic to aggregated interfaces.

To include the bits in the IP address of the IPv4 or IPv6 payload as well as the first label in the hash key, include the **family mpls** statement at the **[edit forwarding-options hash-key]** hierarchy level:

```
[edit forwarding-options hash-key]
family mpls {
  label-1;
  payload {
    ip;
  }
}
```

To include the bits of the IP address of the IPv4 or IPv6 payload as well as both the first label and the second label in the hash key, include the **family mpls** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level.

```
[edit forwarding-options hash-key family mpls]
family mpls {
  no-interface-index;
  label-1;
  label-2;
  label-3;
  no-labels;
  payload {
    ip {
      layer-3-only;
      port-data {
        source-msb;
        source-lsb;
        destination-msb;
        destination-lsb;
      }
    }
  }
}
```

To include only Layer 3 IP information in the hash, specify the **layer-3-only** option.
 To include the most significant byte of the source port, specify the **source-msb** option.
 To include the least significant byte of the source port, specify the **source-lsb** option.
 To include the most significant byte of the destination port, specify the **destination-msb** option. To include the least significant byte of the destination port, specify the **destination-lsb** option.

By default, the most significant bit and least significant bit of the source and destination port fields are hashed. To select specific bits to be hashed, include one or more of the `source-msb`, `source-lsb`, `destination-msb`, and `destination-lsb` options at the `[edit forwarding-options hash-key family mpls payload ip port-data]` hierarchy level. To prevent all four bits from being hashed, include the `layer-3-only` option at the `[edit forwarding-options hash-key family mpls payload ip]` hierarchy level.

In an L2VPN scenario, the router could encounter a reordering complication. When a burst of traffic pushes the customer traffic bandwidth to exceed its limits, the traffic may be affected in mid flow. Packets may be reordered as a result.

You can configure the EXP bit of the top label to be excluded from the hash calculations to avoid the reordering complication. To exclude the EXP bit of the first label from the hash calculations, include the `no-label-1-exp` statement at the `[edit forwarding-options hash-key family mpls]` hierarchy level:

```
[edit forwarding-options hash-key]
family mpls {
  label-1;
  no-label-1-exp;
  payload {
    ip;
  }
}
```

You must also configure the `label-1` statement when configuring the `no-label-1-exp` statement.

Load Balancing Based on Layer 2 MAC Information

To load-balance traffic based on Layer 2 media access control (MAC) information, include the `family multiservice` statement at the `[edit forwarding-options hash-key]` hierarchy level:

```
family multiservice {
  destination-mac;
  source-mac;
}
```

To include the destination-address MAC information in the hash key, include the `destination-mac` option. To include the source-address MAC information in the hash key, include the `source-mac` option.



NOTE: You can configure per-packet load balancing to optimize VPLS traffic flows across multiple paths.



NOTE: This feature is not supported on J-series Services Routers.

Examples: Configuring Per-Packet Load Balancing

Perform per-packet load balancing for all routes:

```
[edit]
policy-options {
  policy-statement load-balancing-policy {
    then {
      load-balance per-packet;
    }
  }
}
routing-options {
  forwarding-table {
    export load-balancing-policy;
  }
}
```

Perform per-packet load balancing only for a limited set of routes:

```
[edit]
policy-options {
  policy-statement load-balancing-policy {
    from {
      route-filter 192.168.10/24 orlonger;
      route-filter 10.114/16 orlonger;
    }
    then {
      load-balance per-packet;
    }
  }
}
routing-options {
  forwarding-table {
    export load-balancing-policy;
  }
}
```


Chapter 7

Summary of Routing Policy Configuration Statements

The following sections explain each of the routing policy configuration statements. The statements are organized alphabetically.

apply-path

Syntax	<code>apply-path path;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> policy-options prefix-list <i>name</i>], [edit policy-options prefix-list <i>name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Expand a prefix list to include all prefixes pointed to by a defined path.
Options	<p><i>path</i>—String of elements composed of identifiers or configuration keywords that points to a set of prefixes. You can include wildcards (enclosed in angle brackets) to match more than one identifier.</p> <p><i>prefix-list name</i>—Name of a list of IP version 4 (IPv4) or IP version 6 (IPv6) prefixes. To create a named list of IP address prefixes, see “Extended Match Conditions Configuration” on page 91.</p>
Usage Guidelines	See “Defining Prefix Lists” on page 111.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

as-path

Syntax	<code>as-path <i>name</i> <i>regular-expression</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define an autonomous system (AS) path regular expression for use in a routing policy match condition.
Options	<p><i>name</i>—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p><i>regular-expression</i>—One or more regular expressions used to match the AS path.</p>
Usage Guidelines	See “Configuring AS Path Regular Expressions” on page 91.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

as-path-group

Syntax	<code>as-path-group <i>group-name</i> { as-path <i>name</i> <i>regular-expression</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define a group containing multiple AS path regular expressions for use in a routing policy match condition.
Options	<p><i>group-name</i>—Name that identifies the AS path group. One or more AS path regular expressions must be listed below the as-path-group hierarchy.</p> <p><i>name</i>—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p><i>regular-expression</i>—One or more regular expressions used to match the AS path.</p>
Usage Guidelines	See “Configuring AS Path Regular Expressions” on page 91.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

community

Syntax community *name* {
 invert-match;
 members [*community-ids*];
 }

Hierarchy Level [edit logical-systems *logical-system-name* policy-options],
 [edit policy-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define a community or extended community for use in a routing policy match condition.

Options *name*—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters. To include spaces in the name, enclose it in quotation marks (" ").

invert-match—Invert the results of the community expression matching.

members *community-ids*—One or more community members. If you specify more than one member, you must enclose all members in brackets.

The format for *community-ids* is:

as-number:community-value

as-number is the AS number and can be a value in the range from 0 through 65,535.
community-value is the community identifier and can be a number in the range from 0 through 65,535.

You also can specify *community-ids* for communities as one of the following well-known community names, which are defined in RFC 1997, *BGP Communities Attribute*:

- **no-export**—Routes containing this community name are not advertised outside a Border Gateway Protocol (BGP) confederation boundary.
- **no-advertise**—Routes containing this community name are not advertised to other BGP peers.
- **no-export-subconfed**—Routes containing this community name are not advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation.

You can explicitly exclude BGP community information with a static route using the **none** option. Include **none** when configuring an individual route in the **route** portion of the **static** statement to override a **community** option specified in the **defaults** portion of the statement.

The format for extended *community-ids* is the following:

type:administrator:assigned-number

type is the type of extended community and can be either a **bandwidth**, **target**, **origin**, **domain-id**, **src-as**, or **rt-import** community or a 16-bit number that identifies a specific BGP extended community. The **target** community identifies the destination to which the route is going. The **origin** community identifies where the route originated. The **domain-id** community identifies the Open Shortest Path First (OSPF) domain from which the route originated. The **src-as** community identifies the autonomous system from which the route originated. The **rt-import** community identifies the route to install in the routing table.



NOTE: For **src-as**, you can specify only an AS number and not an IP address. For **rt-import**, you can specify only an IP address and not an AS number.

administrator is the administrator. It is either an AS number or an IPv4 address prefix, depending on the type of extended community.

assigned-number identifies the local provider.

The format for linking a bandwidth with an AS number is:

bandwidth:as-number:bandwidth

as-number specifies the AS number and **bandwidth** specifies the bandwidth in bytes per second.



NOTE: Beginning with JUNOS Release 9.1, the AS number range has been extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. For extended communities, you can now specify an AS value in plain-number format in the range from 1 through 4,294,967,295. To configure an extended community for a **target** or an **origin** that includes a 4-byte AS number in plain-number format, you must append the letter “L” to the end of the number. For example, to configure a target for the 4-byte AS number 334,324 with an assigned number of 132, include the following statement: **target:334324L:132**.

Beginning with JUNOS Release 9.2, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *< 16-bit high-order value in decimal > . < 16-bit low-order value in decimal > .* For example, the 4-byte AS number of 65546 in plain-number format is represented as 1.10 in the AS-dot notation format. The **target** and **origin** extended communities support the AS-dot notation format for 4-byte AS numbers.

Usage Guidelines

See “Configuring Communities” on page 97, “Inverting Community Matches” on page 109, and “Configuring Link Bandwidth” on page 109.

Required Privilege Level

routing—To view this statement in the configuration.

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

condition

Syntax	condition <i>condition-name</i> { if-route-exists <i>address</i> table <i>table-name</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Define a policy condition based on the existence of routes in specific tables for use in BGP export policies.
Options	if-route-exists <i>address</i> —Specify the address of the route in question. table <i>table-name</i> —Specify a routing table.
Usage Guidelines	See “Configuring the Condition Statement” on page 127.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

damping

Syntax	<pre>damping <i>name</i> { disable; half-life <i>minutes</i>; max-suppress <i>minutes</i>; reuse <i>number</i>; suppress <i>number</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define route flap damping properties to set on BGP routes.
Options	<p>disable—Disable damping on a per-prefix basis. Any damping state that is present in the routing table for a prefix is deleted if damping is disabled.</p> <p>half-life <i>minutes</i>—Decay half-life. <i>minutes</i> is the interval after which the accumulated figure-of-merit value is reduced by half if the route remains stable. Range: 1 through 45 Default: 15 minutes</p> <p>max-suppress <i>minutes</i>—Maximum hold-down time. <i>minutes</i> is the maximum time that a route can be suppressed no matter how unstable it has been. Range: 1 through 720 Default: 60 minutes</p> <p><i>name</i>—Name that identifies the set of damping parameters. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p>reuse <i>number</i>—Reuse threshold. <i>number</i> is the figure-of-merit value below which a suppressed route can be used again. Range: 1 through 20,000 Default: 750 (unitless)</p> <p>suppress <i>number</i>—Cutoff (suppression) threshold. <i>number</i> is the figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. Range: 1 through 20,000 Default: 3000 (unitless)</p>
Usage Guidelines	See “Configuring Flap Damping Parameters” on page 131.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

export

Syntax export [*policy-names*];

Hierarchy Level [edit protocols bgp],
 [edit protocols bgp group *group-name*],
 [edit protocols bgp *group-name* neighbor *address*],
 [edit protocols bgp group *group-name* neighbor *address* out-delay *seconds*],
 [edit protocols bgp out-delay *seconds*],
 [edit protocols dvmrp],
 [edit protocols isis],
 [edit protocols ldp],
 [edit protocols msdp],
 [edit protocols msdp group *group-name*],
 [edit protocols msdp group *group-name* peer *peer-address*],
 [edit protocols msdp peer *address*],
 [edit protocols ospf],
 [edit protocols ospf3],
 [edit protocols pim rp bootstrap family (inet | inet6)],
 [edit rip group *group-name*],
 [edit ripng group *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* protocols bgp group *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor
 address],
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor
 address out-delay *seconds*],
 [edit logical-systems *logical-system-name* protocols dvmrp],
 [edit logical-systems *logical-system-name* protocols isis],
 [edit logical-systems *logical-system-name* protocols ldp],
 [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 ospf],
 [edit logical-systems *logical-system-name* protocols ospf3],
 [edit logical-systems *logical-system-name* protocols pim rp bootstrap family (inet | inet6)],
 [edit logical-systems *logical-system-name* rip group *group-name*],
 [edit logical-systems *logical-system-name* ripng group *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor
 address],
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor
 address out-delay *seconds*],
 [edit routing-instances *routing-instance-name* protocols bgp out-delay *seconds*],
 [edit routing-instances *routing-instance-name* protocols dvmrp],
 [edit routing-instances *routing-instance-name* protocols isis],
 [edit routing-instances *routing-instance-name* protocols ldp],
 [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],
[edit routing-instances routing-instance-name protocols ospf],
[edit routing-instances routing-instance-name protocols ospf3],
[edit routing-instances routing-instance-name protocols pim rp bootstrap family (inet |
  inet6)],
[edit routing-instances routing-instance-name protocols rip group group-name],
[edit routing-instances routing-instance-name protocols ripng group group-name]
```

Release Information	Statement introduced before JUNOS Release 7.4.
Description	Apply one or more policies to routes being exported from the routing table into a routing protocol.
Options	<i>policy-names</i> —Names of one or more policies defined with a <i>policy-statement</i> statement.
Usage Guidelines	See “Applying a Routing Policy” on page 58.
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 protocols bgp],
 [edit logical-systems *logical-system-name* protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp],
 [edit protocols bgp group *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp group *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor
 address],
 [edit protocols bgp group *group-name* neighbor *address*],
 [edit protocols dvmrp],
 [edit logical-systems *logical-system-name* protocols dvmrp],
 [edit routing-instances *routing-instance-name* protocols dvmrp],
 [edit protocols ldp],
 [edit logical-systems *logical-system-name* protocols ldp],
 [edit routing-instances *routing-instance-name* protocols ldp],
 [edit protocols msdp],
 [edit logical-systems *logical-system-name* protocols msdp],
 [edit routing-instances *routing-instance-name* protocols msdp],
 [edit protocols msdp peer *address*],
 [edit logical-systems *logical-system-name* protocols msdp peer *address*],
 [edit routing-instances *routing-instance-name* protocols msdp peer *address*],
 [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 protocols msdp group *group-name* peer *address*],
 [edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
 [edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer
 address],
 [edit protocols ospf],
 [edit logical-systems *logical-system-name* protocols ospf],
 [edit routing-instances *routing-instance-name* protocols ospf],
 [edit protocols ospf3],
 [edit logical-systems *logical-system-name* protocols ospf3],
 [edit routing-instances *routing-instance-name* protocols ospf3],
 [edit protocols pim],
 [edit logical-systems *logical-system-name* protocols pim],
 [edit routing-instances *routing-instance-name* protocols pim],
 [edit protocols pim rp bootstrap family (inet | inet6)],
 [edit logical-systems *logical-system-name* protocols pim rp bootstrap family (inet | inet6)],
 [edit routing-instances *routing-instance-name* protocols pim rp bootstrap family (inet |
 inet6)],
 [edit protocols rip],
 [edit logical-systems *logical-system-name* protocols rip],
 [edit routing-instances *routing-instance-name* protocols rip],
 [edit protocols rip group *group-name*],
 [edit logical-systems *logical-system-name* protocols rip group *group-name*],
 [edit routing-instances *routing-instance-name* protocols rip group *group-name*],
 [edit protocols rip group *group-name* neighbor *address*],

```
[edit logical-systems logical-system-name protocols rip group group-name neighbor
  address],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor
  address],
[edit protocols ripng],
[edit logical-systems logical-system-name protocols ripng],
[edit routing-instances routing-instance-name protocols ripng],
[edit protocols ripng group group-name],
[edit logical-systems logical-system-name protocols ripng group group-name],
[edit routing-instances routing-instance-name protocols ripng group group-name],
[edit protocols ripng group group-name neighbor address],
[edit logical-systems logical-system-name protocols ripng group group-name neighbor
  address],
[edit routing-instances routing-instance-name protocols ripng group group-name neighbor
  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 a routing protocol.
Options	<i>policy-names</i> —Names of one or more policies defined with a <i>policy-statement</i> statement.
Usage Guidelines	See “Applying a Routing Policy” on page 58.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

policy-options

Syntax	<code>policy-options { ... }</code>
Hierarchy Level	[edit]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure routing policy.
Options	The statements are explained separately.
Usage Guidelines	See “Defining Routing Policies” on page 39.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

policy-statement

Syntax

```

policy-statement policy-name {
  term term-name {
    from {
      family family-name;
      match-conditions;
      policy subroutine-policy-name;
      prefix-list prefix-list-name;
      prefix-list-filter prefix-list-name match-type <actions>;
      route-filter destination-prefix match-type <actions>;
      source-address-filter source-prefix match-type <actions>;
    }
    to {
      match-conditions;
      policy subroutine-policy-name;
    }
    then actions;
  }
}

```

Hierarchy Level [edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define a routing policy, including subroutine policies.

Options *actions*—(Optional) One or more actions to take if the conditions match. The actions are described in Table 15 on page 48 and Table 16 on page 49.

family family-name—(Optional) Specify an address family protocol. Specify *inet* for an IPv4 address protocol. Specify *inet6* for a 128-bit IPv6 address protocol, and to enable interpretation of IPv6 router filter addresses. When *family* is not specified, the router uses the default IPv4 setting.

from—(Optional) Match a route based on its source address.

match-conditions—(Optional in *from* statement; required in *to* statement) One or more conditions to use to make a match. The qualifiers are described in Table 14 on page 42.

policy subroutine-policy-name —Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”). For information about how to configure subroutines, see “Configuring Subroutines” on page 123.

policy-name—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

prefix-list *prefix-list-name* —Name of a list of IPv4 or IPv6 prefixes. To create a named list of IP address prefixes, see “Extended Match Conditions Configuration” on page 91.

prefix-list-filter *prefix-list-name*—Name of a prefix list to evaluate using qualifiers; *match-type* is the type of match (see Table 25 on page 115), and *actions* is the action to take if the prefixes match.

route-filter *prefix match-type <actions>*—(Optional) List of routes on which to perform an immediate match; *source-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see Table 24 on page 113), and *actions* is the action to take if the *source-prefix* matches.

source-address-filter *source-prefix match-type <actions>*—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. *source-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see Table 25 on page 115), and *actions* is the action to take if the *source-prefix* matches.

term *term-name*—Name that identifies the term.

to—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

then—(Optional) Actions to take on matching routes. The actions are described in Table 15 on page 48 and Table 16 on page 49.

Usage Guidelines See “Defining Routing Policies” on page 39 and “Extended Match Conditions Configuration” on page 91.

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

prefix-list

Syntax	prefix-list <i>name</i> { <i>ip-addresses</i> ; apply-path <i>path</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define a list of IPv4 or IPv6 address prefixes for use in a routing policy statement or firewall filter statement.
Options	<p><i>name</i>—Name that identifies the list of IPv4 or IPv6 address prefixes.</p> <p><i>ip-addresses</i>—List of IPv4 or IPv6 address prefixes, one IP address per line in the configuration.</p> <p>The remaining statement is explained separately in this chapter.</p>
Usage Guidelines	See “Configuring Prefix Lists” on page 110.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

prefix-list-filter

Syntax	prefix-list-filter <i>prefix-list-name</i> <i>match-type</i> < <i>actions</i> >;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Evaluate a list of prefixes within a prefix list using specified qualifiers.
Options	<p><i>prefix-list-name</i>—Name of the prefix list to evaluate.</p> <p><i>match-type</i>—Prefix length qualifiers.</p> <p><<i>actions</i>>—(Optional) Actions to take on match.</p>
Usage Guidelines	See “Configuring Prefix Lists” on page 110.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

Part 3

Firewall Filters

- Firewall Filter Overview on page 159
- Firewall Filter Configuration on page 163
- Policer Overview on page 231
- Policer Configuration on page 233
- Summary of Firewall Filter and Policer Configuration Statements on page 257

Chapter 8

Firewall Filter Overview

Firewall filters allow you to filter packets based on their components and to perform an action on packets that match the filter.



NOTE: The JUNOS software provides a *policy framework*, which is a collection of JUNOS policies that include routing policies and firewall filter policies. These policies share some fundamental similarities. (For information about the similarities and differences among these policies, see “Policy Framework Overview” on page 3.) However, when you refer to a firewall filter policy in the firewall filters part of the manual, the term *firewall filter* is used.

Depending on the hardware configuration of the routing platform, you can use firewall filters for the following purposes:

- On routing platforms equipped with an Internet Processor II application-specific integrated circuit (ASIC), you can control *data packets*, which are chunks of data transiting the routing platform as they are forwarded from a source to a destination.
- On all routing platforms, you can control the *local packets*, which are chunks of data that are destined for or sent by the Routing Engine.

With the Internet Processor II ASIC, you can use filters on data packets passing through the routing platform to provide protocol-based firewalls, thwart denial-of-service (DoS) attacks, prevent falsifying of source addresses, create access control lists, and implement rate limiting (policing). (Use the `show chassis hardware` command to determine whether a routing platform has an Internet Processor or an Internet Processor II ASIC.)

You can use the filters to restrict the local packets that pass from the routing platform's physical interfaces to the Routing Engine. Such filters are useful in protecting the IP services that run on the Routing Engine, such as telnet, Secure Shell (ssh), and the Border Gateway Protocol (BGP), from denial-of-service attacks. You can define input filters, which affect only inbound traffic destined for the Routing Engine, and output filters, which affect only outbound traffic sent from the Routing Engine. You can also use policing, or rate limiting, to provide a finer level of control over local packets destined for the Routing Engine.



NOTE: In the remainder of the firewall filters part of this manual, the term *packets* refers to both data and local packets unless explicitly stated otherwise.

You can apply firewall filters to packets entering or leaving the routing platform on one, more than one, or all interfaces. For each interface, you can apply a firewall filter to incoming or outgoing traffic, or both, and the same filter can be used for both.

You can define firewall filters that apply to IP version 4 (IPv4), IP version 6 (IPv6), or Multiprotocol Label Switching (MPLS) traffic.

There is no limit to the number of filters and counters you can set, but there are some practical considerations. More counters require more terms, and a large number of terms can take a long time to process during a commit. However, filters with more than 1000 terms and counters have been implemented successfully.

This chapter describes the following topics:

- Firewall Filter Components on page 160
- Supported Standards on page 161

Firewall Filter Components

In a firewall filter, you first define the address structure type (IPv4, IPv6, or MPLS) and then you define one or more terms that specify the filtering criteria and the action to take if a match occurs. Each term consists of two components:

- Match conditions—Values or fields that the packet must contain. You can define various match conditions, including the IP source address field, IP destination address field, Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port field, IP protocol field, Internet Control Message Protocol (ICMP) packet type, IP options, TCP flags, incoming logical or physical interface, and outgoing logical or physical interface.
- Action—Specifies what to do if a packet matches the match conditions. Possible actions are to accept, discard, or reject a packet, go to the next term, or take no action. In addition, statistical information can be recorded for a packet: it can be counted, logged, or sampled.

The order of the terms within a firewall filter is significant. Packets are tested against each term in the order in which it is listed in the configuration. When the first matching conditions are found, the action associated with that term is applied to the packet and the evaluation of the firewall filter ends, unless the **next term** action modifier is included. If the **next term** action is included, the matching packet is then evaluated against the next term in the firewall filter; otherwise, the matching packet is not evaluated against subsequent terms in the firewall filter.

If, after all terms are evaluated, a packet matches no terms in a filter, the packet is silently discarded.

If a packet arrives on an interface and a firewall filter is not configured for the incoming traffic on that interface, the packet is accepted by default.

Although policing, traffic sampling, and forwarding are configured as firewall filters, they are documented in separate parts of this manual. For information about policing, see “Policer Configuration” on page 233. For information about traffic sampling and forwarding, see “Traffic Sampling and Forwarding Configuration” on page 279.

Supported Standards

The JUNOS software supports the following RFCs related to filtering:

- RFC 792, *Internet Control Message Protocol (ICMP)*
- RFC 2373, *IP Version 6 Addressing Architecture*
- RFC 2460, *Internet Protocol, Version 6 (IPv6)*
- RFC 2474, *Definition of the Differentiated Services (DS) Field*
- RFC 2475, *An Architecture for Differentiated Services*
- RFC 2597, *Assured Forwarding PHB*
- RFC 2598, *An Expedited Forwarding PHB*

Chapter 9

Firewall Filter Configuration

To configure firewall filters, you include statements at the `[edit firewall]` hierarchy level of the configuration:

```
[edit firewall]
family family-name {
  filter filter-name {
    accounting-profile name;
    interface-specific;
    term term-name {
      filter filter-name;
      from {
        match-conditions;
      }
      then {
        action;
        action-modifiers;
      }
    }
  }
}
service-filter filter-name {
  term term-name {
    from {
      match-conditions;
    }
    then {
      action;
      action-modifiers;
    }
  }
}
simple-filter filter-name {
  term term-name {
    from {
      match-conditions;
    }
    then {
      action;
      action-modifiers;
    }
  }
}
}
```



NOTE: Stateless firewall filtering is not supported on the interfaces you configure as tunnel sources. The firewall filter affects only the transit packets exiting the tunnel.

This chapter describes the following tasks for configuring firewall filters:

- Minimum Firewall Filter Configuration on page 164
- Configuring Firewall Filters on page 165
- Configuring Service Filters on page 203
- Configuring Simple Filters on page 203
- Applying Firewall Filters to Interfaces on page 205
- Configuring Firewall Filters for Logical Systems on page 208
- Configuring Accounting on page 221
- Configuring Filter-Based Forwarding on page 222
- Configuring Forwarding Table Filters on page 225
- Configuring Firewall Filter System Logging on page 227

Minimum Firewall Filter Configuration

To configure a firewall filter, you must perform at least the following tasks:

- Configure firewall filters—To configure firewall filters, include the **family** *family-name* statement and one or more **filter** statements at the [edit firewall] hierarchy level:

```
[edit firewall]
family family-name {
  filter filter-name {
    term term-name {
      from {
        match-conditions;
      }
      then {
        action;
        action-modifiers;
      }
    }
  }
}
```

- Apply firewall filters to interfaces—Firewall filters control local packets to and from the Routing Engine if they are applied to the loopback interface, **lo0**. With the Internet Processor II application-specific integrated circuit (ASIC), firewall filters can control data packets through the routing platform when they are applied to an external interface. To have a firewall filter take effect, you must apply it to an interface by including the **filter** statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family-name*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family family-name]
```

```
filter {
  input filter-name;
  output filter-name;
}
```

Configuring Firewall Filters

To configure firewall filters, include the **firewall** statement at the [edit] hierarchy level:

```
[edit]
firewall {
  family family-name {
    filter filter-name {
      accounting-profile name;
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
}
```



NOTE: You can specify IP version 4 (IPv4) filters at either the [edit firewall] hierarchy level or the [edit firewall family inet] hierarchy level.

The following sections describe the components of the **firewall** statement and provide examples of configuring firewall filters:

- Configuring the Family Address Type on page 166
- Configuring the Filter Name on page 166
- Configuring the Filter Terms on page 166
- Configuring a Filter Match Statement on page 167
- Configuring a Filter Action Statement on page 168
- How Firewall Filters Are Evaluated on page 173
- Filter Match Conditions on page 174
- How Firewall Filters Test a Packet's Protocol on page 191
- Configuring a Filter Within a Filter on page 193
- Examples: Defining Firewall Filters on page 194

Configuring the Family Address Type

To configure the family address type for a firewall filter, include the **family** statement at the [edit firewall] hierarchy level:

```
[edit firewall]
family { ... }
```

Specify **inet** to filter IPv4 packets. Specify **inet6** to filter IP version 6 (IPv6) packets. Specify **mpls** to filter Multiprotocol Label Switching (MPLS) packets. Specify **vpls** to filter virtual private LAN service (VPLS) packets. Specify **ccc** to filter Layer 2 switching cross-connects.

On the MX-series router only, to filter Layer 2 packets in a bridging environment, specify **bridge**.

Specify **any** to filter packets based upon protocol-independent fields. This type of filter allows the firewall to match on the following protocol-independent match conditions: **forwarding-class**, **forwarding-class-except**, **interface**, **packet-length**, **packet-length-except**, **interface-index**, and **interface-set**.



NOTE: The address type family **any** is supported on the loopback (lo0) interface only.

Configuring the Filter Name

To configure the filter name, include the **filter** statement:

```
filter filter-name { ... }
```

For IPv4 traffic, configure the filter name at the [edit firewall family inet] hierarchy level. For IPv6 traffic, configure the filter name at the [edit firewall family inet6] hierarchy level. The filter name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").

Configuring the Filter Terms

Each firewall filter consists of one or more *terms*. To configure a term, include the **term** statement:

```
term term-name { ... }
```

For IPv4 traffic, configure the filter terms at the [edit firewall family inet filter filter-name] hierarchy level. For IPv6 traffic, configure the filter terms at the [edit firewall family inet6 filter filter-name] hierarchy level. For MPLS traffic, configure the filter terms at the [edit firewall family mpls filter filter-name] hierarchy level. For Layer 2 bridging traffic, configure the filter terms at the [edit firewall family bridge filter] hierarchy level.



NOTE: Layer 2 bridging is supported only on the MX-series routers. For more information about how to configure Layer 2 bridging, see the *JUNOS Network Interfaces Configuration Guide*, the *JUNOS MX-series Layer 2 Configuration Guide*, and the *JUNOS Feature Guide*.

The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").

Each term name must be unique within a filter.

You can specify multiple terms in a filter, effectively chaining together a series of match-action operations to apply to the packets on an interface. You can also use the **next term** action so that, when a match condition is met, the evaluation continues to the next term, rather than terminating.

Firewall filter terms are evaluated in the order in which you specify them in the configuration. To reorder terms, use the configuration mode **insert** command. For example, the command **insert term up before term start** places the term **up** before the term **start**.

Configuring a Filter Match Statement

In a firewall filter term, you can define conditions used to match the components of a packet. To configure match conditions, include the **from** statement:

```
from {
    match-conditions;
}
```

For IPv4 traffic, configure the match conditions at the **[edit firewall family inet filter filter-name term term-name]** hierarchy level. For IPv6 traffic, configure the match conditions at the **[edit firewall family inet6 filter filter-name term term-name]** hierarchy level. For MPLS traffic, configure the filter terms at the **[edit firewall family mpls filter filter-name term term-name]** hierarchy level.

For traffic in a bridging environment, configure the match conditions at the **[edit firewall family bridge filter filter-name term term-name]** hierarchy level. The **bridge** option is supported only on MX-series routers. You can configure filters to match on Layer 2, Layer 3, and Layer 4 fields. Although interfaces configured for bridging can receive both IPv4 and IPv6 traffic, Layer 3 and Layer 4 match conditions currently support IPv4 traffic only. For more information about the filter match conditions supported for traffic in a bridging environment, see Table 31 on page 182.

You can specify zero or more match conditions in a single **from** statement. For a match to occur, the packet must match all the conditions in the term. For more information about match conditions, see “Filter Match Conditions” on page 174.

The **from** statement is optional. If you omit it, all packets are considered to match.

Configuring a Filter Action Statement

In a firewall filter term, you can specify the action to take if the packet matches the conditions you have configured in the term. To configure a filter action, include the **then** statement:

```
then {
    action;
    action-modifiers;
}
```

If you omit the **then** statement or do not specify an action, the packets that match the conditions in the **from** statement are accepted.

For IPv4 traffic, configure the filter action at the [edit firewall family inet filter *filter-name* term *term-name*] hierarchy level. For IPv6 traffic, configure the filter action at the [edit firewall family inet6 filter *filter-name* term *term-name*] hierarchy level. For MPLS traffic, configure the filter action at the [edit firewall family mpls filter *filter-name* term *term-name*] hierarchy level.

For Layer 2 traffic in a bridging environment, configure the filter action at the [edit firewall family bridge filter *filter-name* term *term-name*] hierarchy level. The **bridge** option is supported only on MX-series routers.



NOTE: We strongly recommend that you always explicitly configure an action in the **then** statement.

You can specify one of the following filter actions:

- **accept**—Accept the packet, which is then sent to its destination.
- **discard**—Discard the packet, which is not processed further.
- **next term**—Evaluate the next term in the firewall filter.
- **reject**—Reject the packet and send a rejection message. Rejected packets can be logged or sampled.
- **routing-instance**—Accept the packet, which is then routed by the specified routing instance. For more information, see “Configuring Filter-Based Forwarding” on page 222.
- **topology**—Accept the packet, which is then installed in the routing table of the specified topology. For more information, see the “Configuring Multitopology Routing” chapter in the *JUNOS Routing Protocols Configuration Guide*.

In the filter **action** statement, you can also specify one or more of the following action modifiers:

- **count**—Add the packet to a count total.
- **forwarding-class**—Specify the packet forwarding class name.
- **log**—Store the packet header information in a buffer within the Packet Forwarding Engine.

- **loss-priority**—Set the packet loss priority (PLP) to **low**, **medium**, or **high**.



NOTE: You must configure tricolor marking policer to set PLP to **medium**.

- **policer**—Apply rate-limiting procedures to the traffic. For more information, see “Policer Configuration” on page 233.
- **sample**—Sample the packet traffic. Apply this option only if you have enabled traffic sampling. For more information, see “Traffic Sampling and Forwarding Configuration” on page 279.
- **syslog**—Store the packet header information on the Routing Engine and log it to the system log.



NOTE: The firewall filter **syslog** action stops logging at a high traffic rate to protect the Routing Engine from an excessive flow of messages.

- **ipsec-sa *sa-name***—Specify an IP Security (IPsec) security association (SA) for the packet. This is used with the **source-address** and **destination-address** match conditions.

You can specify only one filter **action** statement (or omit it), but you can specify any combination of action modifiers. For the action or action modifier to take effect, all conditions in the **from** statement must match. If you specify **log** as one of the actions in a term, this constitutes a termination action; whether any additional terms in the filter are processed depends on the traffic through the filter.

The action modifier operations carry a default **accept** action. For example, if you specify an action modifier and do not specify an action, the specified action modifier is implemented and the packet is accepted.



NOTE: You cannot configure both the **loss-priority** and **three-color-policer** action modifiers for the same firewall filter term.

Policing uses a specific type of action, known as a policer action. For more information, see “Policer Configuration” on page 233.

For more information about forwarding classes and loss priority, see the *JUNOS Class of Service Configuration Guide*.

Table 28 on page 169 shows the complete list of filter actions and action modifiers.

Table 28: Firewall Filter Actions and Action Modifiers

Action or Action Modifier	Description
Actions	

Table 28: Firewall Filter Actions and Action Modifiers (*continued*)

Action or Action Modifier	Description
accept	Accept a packet.
count <i>counter-name</i>	Count the packet in the specified counter.
dscp	Set the IPv4 or the IPv6 Differentiated Services code point (DSCP) bit to 0.
discard	Discard a packet silently, without sending an Internet Control Message Protocol (ICMP) message. Discarded packets are available for logging and sampling.
forwarding-class <i>class</i>	Classify the packet into one of the following forwarding classes: as , assured-forwarding , best-effort , expedited-forwarding , or network-control .
ipsec-sa <i>ipsec-sa</i>	Use the specified IPsec security association.
load-balance <i>group-name</i>	Use the specified load-balancing group.
logical-system <i>logical-system-name</i>	Use the specified logical system. This action is supported for both IPv4 and IPv6 firewall filters.
loss-priority (high low medium)	Set the loss priority level for packets.
out-of-profile	Indicate that the upper or lower bound of a policer has been met and starvation of queues is possible. The packets are marked as out of the profile of the policer. This action is supported on the J-series Services Router only as part of strict priority queuing. Out-of-profile packets are queued only if the port is not congested.
next <i>term</i>	Continue to the next term for evaluation.
next-hop-group <i>group-name</i>	Use the specified next-hop group.
policer <i>policer-name</i>	Rate-limit packets based on the specified policer.
port-mirror	Port-mirror the packets.
prefix-action <i>name</i>	Count or police packets based on the specified action name.
reject <i>message-type</i>	Discard a packet, sending an ICMPv4 or an ICMPv6 destination unreachable message. Rejected packets can be logged or sampled if you configure either the sample or the syslog action modifier. You can specify one of the following message codes: administratively-prohibited (default), bad-host-tos , bad-network-tos , host-prohibited , host-unknown , host-unreachable , network-prohibited , network-unknown , network-unreachable , port-unreachable , precedence-cutoff , precedence-violation , protocol-unreachable , source-host-isolated , source-route-failed , or tcp-reset . If you specify tcp-reset , a Transmission Control Protocol (TCP) reset is returned if the packet is a TCP packet. Otherwise, nothing is returned.
routing-instance <i>routing-instance</i>	Specify a routing instance to which packets are forwarded.
sample	Sample the packets.

Table 28: Firewall Filter Actions and Action Modifiers (*continued*)

Action or Action Modifier	Description
topology <i>topology-name</i>	Specify a topology to which packets are forwarded.
Action Modifiers	
count <i>counter-name</i>	Number of packets passing this filter/term/policer. The name can contain letters, numbers, underscores (_), and hyphens (-), and can be up to 64 characters long. A counter name is specific to the filter that uses it, so all interfaces that use the same filter increment the same counter.
forwarding-class <i>class-name</i>	Particular forwarding class.
ipsec-sa <i>sa-name</i>	IPsec SA for the packet. Used with the source-address and destination-address match conditions.
log	Log the packet header information in a buffer within the Packet Forwarding Engine. You can access this information by issuing the show firewall log command at the command-line interface (CLI).
loss-priority <i>priority</i>	Set the PLP to low or high . You cannot also configure the three-color-policer action modifier for the same firewall filter term. These two action modifiers are mutually exclusive.
policer <i>policer-name</i>	Apply rate limits to the traffic using the named policer.
sample	Sample the traffic on the interface. Use this modifier only when traffic sampling is enabled. For more information, see “Traffic Sampling and Forwarding Configuration” on page 279.
syslog	Store the packet header information on the Routing Engine and log it to the system log.
three-color-policer <i>policer-name</i>	Apply rate limits to the traffic using the tricolor marking policer. You cannot also configure the loss-priority action modifier for the same firewall filter term. These two action modifiers are mutually exclusive.

Example: Configure a Filter Action Statement

Count, sample, and accept the traffic:

```
term all {
  then {
    count sam-1;
    sample; # default action is accept
  }
}
```

Display the packet counter:

```
user@host> show firewall filter sam
Filter:
Counters:
```

Name	Bytes	Packets
sam		
sam-1	98	8028

Display the firewall log output:

```
user@host> show firewall log
```

Time	Filter	A	Interface	Pro	Source address	Destination address
23:09:09	-	A	at-2/0/0.301	TCP	10.2.0.25	10.211.211.1:80
23:09:07	-	A	at-2/0/0.301	TCP	10.2.0.25	10.211.211.1:56
23:09:07	-	A	at-2/0/0.301	ICM	10.2.0.25	10.211.211.1:49552
23:02:27	-	A	at-2/0/0.301	TCP	10.2.0.25	10.211.211.1:56
23:02:25	-	A	at-2/0/0.301	TCP	10.2.0.25	10.211.211.1:80
23:01:22	-	A	at-2/0/0.301	ICM	10.2.2.101	10.211.211.1:23251
23:01:21	-	A	at-2/0/0.301	ICM	10.2.2.101	10.211.211.1:16557
23:01:20	-	A	at-2/0/0.301	ICM	10.2.2.101	10.211.211.1:29471
23:01:19	-	A	at-2/0/0.301	ICM	10.2.2.101	10.211.211.1:26873

This output file contains the following fields:

- **Time**—Time at which the packet was received (not shown in the default).
- **Filter**—Name of a filter that has been configured with the `filter` statement at the `[edit firewall]` hierarchy level. A hyphen (-) or the abbreviation `pfe` indicates that the packet was handled by the Packet Forwarding Engine. A space (no hyphen) indicates that the packet was handled by the Routing Engine.
- **A**—Filter action:
 - A—Accept (or next term)
 - D—Discard
 - R—Reject
- **Interface**—Interface on which the filter is configured.



NOTE: We strongly recommend that you always explicitly configure an action in the `then` statement.

- **Pro**—Packet's protocol name or number.
- **Source address**—Source IP address in the packet.
- **Destination address**—Destination IP address in the packet.

Display the sampling output:

```
user@host> show log /var/tmp/sam
```

```
# Apr 7 15:48:50
```

Time	Dest	Src	Dest	Src	Proto	TOS	Pkt	Intf	IP	TCP
	addr	addr	port	port			len	num	frag	flags
Apr 7 15:48:54	192.168.9.194	192.168.9.195	0	0	1	0x0	84	8	0x0	0x0
Apr 7 15:48:55	192.168.9.194	192.168.9.195	0	0	1	0x0	84	8	0x0	0x0
Apr 7 15:48:56	192.168.9.194	192.168.9.195	0	0	1	0x0	84	8	0x0	0x0



NOTE: When you enable reverse path forwarding (RPF) on an interface with an input filter for firewall log and count, the input firewall filter does not log the packets rejected by RPF, although the rejected packets are counted. To log the rejected packets, use an RPF check fail filter.

For more information about sampling output, see “Configuring a Forwarding Table Filter” on page 286.

Example: Set the DSCP Bit to 0

Set the DSCP bit to 0 using a firewall filter:

```
firewall {
  filter filter1 {
    term 1 {
      from {
        dscp 2;
      }
      then {
        dscp 0;
        forwarding-class best-effort;
      }
    }
    term 2 {
      from {
        dscp 3;
      }
      then {
        forwarding-class best-effort;
      }
    }
  }
}
```

Apply this filter to the logical interface corresponding to the VPN routing and forwarding (VRF) instance:

```
interfaces so-0/1/0 {
  unit 0 {
    family inet {
      filter input filter1;
    }
  }
}
```

How Firewall Filters Are Evaluated

When a firewall filter consists of a single term, the filter is evaluated as follows:

- If the packet matches all the conditions, the action in the **then** statement is taken.

- If the packet matches all the conditions, and if there is no action specified in the **then** statement, the default action **accept** is used.
- If the packet does not match all the conditions, it is discarded.

When a firewall filter consists of more than one term, the filter is evaluated sequentially:

- The packet is evaluated against the conditions in the **from** statement in the first term.
- If the packet matches, the action in the **then** statement is taken and, if the **next term** action is not used, the evaluation ends. Subsequent terms in the filter are not evaluated.
- If the packet matches, the action in the **then** statement is taken; if the **next term** action is present, the evaluation continues to the next term.
- If the packet does not match, it is evaluated against the conditions in the **from** statement in the second term.

This process continues until either the packet matches the **from** conditions in one of the subsequent terms or there are no more terms.

- If a packet passes through all the terms in the filter without matching any of them, it is discarded.

If a term does not contain a **from** statement, the packet is considered to match and the action in the term's **then** statement is taken.

If a term does not contain a **then** statement or if you do not configure an action in the **then** statement, and if the packet matches the conditions in the term's **from** statement, the packet is accepted.

Each firewall filter has an implicit discard action at the end of the filter, which is equivalent to the following explicit filter term:

```
term implicit-rule {
  then discard;
}
```

Therefore, if a packet matches none of the terms in the filter, it is discarded.

Filter Match Conditions

In the **from** statement in the **firewall filter** term, you specify conditions that the packet must match for the action in the **then** statement to be taken. All conditions in the **from** statement must match for the action to be taken. The order in which you specify match conditions is not important, because a packet must match all the conditions in a term for a match to occur.

If you specify no match conditions in a term, that term matches all packets.

An individual condition in a **from** statement can contain a list of values. For example, you can specify numeric ranges or multiple source or destination addresses. When a condition defines a list of values, a match occurs if one of the values in the list matches the packet.

Individual conditions in a **from** statement can be negated. When you negate a condition, you are defining an explicit mismatch. If a packet matches a negated condition, it is immediately considered not to match the **from** statement, and the next term in the filter is evaluated, if there is one; if there are no more terms, the packet is discarded.

Match conditions are grouped into categories depending upon how you specify the condition. You can specify the following conditions:

- Specifying Numeric Range Filter Match Conditions on page 175
- Specifying Address Filter Match Conditions on page 184
- Specifying Bit-Field Filter Match Conditions on page 188
- Specifying Class-Based Filter Match Conditions on page 190
- Filtering Smaller Packets on page 191

Specifying Numeric Range Filter Match Conditions

Numeric range filter conditions match packet fields that can be identified by a numeric value, such as port and protocol numbers. For numeric range filter match conditions, you specify a keyword that identifies the condition and a single value or a range of values that a field in a packet must match. Table 29 on page 176 describes the numeric range filter match conditions for IPv4 addresses, and Table 30 on page 179 describes them for IPv6 addresses.

Table 31 on page 182 describes the numeric range filter match conditions for Layer 2 bridging traffic on MX-series routers. For more information about how to configure Layer 2 services on the MX-series routers, see the *JUNOS Network Interfaces Configuration Guide*, the *JUNOS Layer 2 Configuration Guide*, and the *JUNOS Feature Guide*.

You can specify the numeric range value in one of the following ways:

- Single number. A match occurs if the value of the field matches the number. For example:

```
source-port 25;
```

- Range of numbers. A match occurs if the value of the field falls within the specified range. The following example matches source ports 1024 through 65,535, inclusive:

```
source-port 1024-65535;
```

- Text synonym for a single number. A match occurs if the value of the field matches the number that corresponds to the synonym. For example:

```
source-port smtp;
```

To specify multiple values in a single match condition, group the values within square brackets following the keyword. For example:

```
source-port [ smtp ftp-data 25 1024-65535 ];
```

To exclude a numeric value, append the string **-except** to the match keyword. For example, the following condition would match only if the source port is not 25:

```
source-port-except 25;
```

The following condition would match only if the port number is not one of those in the list:

```
source-port-except [ smtp ftp-data 666 1024-65535 ];
```



NOTE: To match only on a source address, destination address, source port or destination port, include the appropriate matching condition (**source-address**, **destination-address**, **source-port**, or **destination-port**, respectively) at the [edit firewall filter *filter-name* term *term-name* from] hierarchy level instead of using the **port** or **address** matching condition at the same hierarchy level.

Table 29: Numeric Range IPv4 Firewall Filter Match Conditions

Match Condition	Description
<i>keyword-except</i>	Negate a match. For example, <i>destination-port-except number</i> .
<i>ah-spi spi-value</i>	IPsec authentication header (AH) security parameter index (SPI) value. Match on this specific SPI value.
<i>ah-spi-except spi-value</i>	IPsec AH SPI value. Do not match on this specific SPI value.
<i>destination-address address</i>	Destination prefix.
<i>destination-mac-address address</i>	Destination media access control (MAC) address of a VPLS packet.

Table 29: Numeric Range IPv4 Firewall Filter Match Conditions (*continued*)

Match Condition	Description
<code>destination-port number</code>	<p>TCP or User Datagram Protocol (UDP) destination port field. You cannot specify both the port and destination-port match conditions in the same term.</p> <p>Normally, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. For more information, see “How Firewall Filters Test a Packet’s Protocol” on page 191.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), xmcp (177).</p>
<code>destination-prefix-list name</code>	Match on the destination prefixes in the specified list name. Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.
<code>dscp number</code>	<p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant 6 bits of this byte form the DSCP. For more information, see the <i>JUNOS Class of Service Configuration Guide</i>.</p> <p>You can specify DSCP in either hexadecimal, binary, or decimal form.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> ■ RFC 2598, <i>An Expedited Forwarding PHB</i>, defines one code point: ef (46). ■ RFC 2597, <i>Assured Forwarding PHB</i>, defines 4 classes, with 3 drop precedences in each class, for a total of 12 code points: <p>af11 (10), af12 (12), af13 (14),</p> <p>af21 (18), af22 (20), af23 (22),</p> <p>af31 (26), af32 (28), af33 (30),</p> <p>af41 (34), af42 (36), af43 (38)</p>
<code>ether-type value</code>	Match on the Ethernet type field of a VPLS packet.
<code>ether-type-except value</code>	Do not match on the Ethernet type field of a VPLS packet.
<code>esp-spi spi-value</code>	IPsec encapsulating security payload (ESP) SPI value. Match on this specific SPI value. You can specify the ESP SPI value in hexadecimal, binary, or decimal form.
<code>esp-spi-except spi-value</code>	IPsec ESP SPI value. Do not match on this specific SPI value.
<code>forwarding-class class</code>	Match on the forwarding class. Specify assured-forwarding , best-effort , expedited-forwarding , or network-control .

Table 29: Numeric Range IPv4 Firewall Filter Match Conditions (*continued*)

Match Condition	Description
<code>forwarding-class-except class</code>	Do not match on the forwarding class. Specify <code>assured-forwarding</code> , <code>best-effort</code> , <code>expedited-forwarding</code> , or <code>network-control</code> .
<code>fragment-offset number</code>	Fragment offset field.
<code>icmp-code number</code>	<p>ICMP code field. This value or keyword provides more specific information than <code>icmp-type</code>. Because the value's meaning depends upon the associated <code>icmp-type</code>, you must specify <code>icmp-type</code> along with <code>icmp-code</code>. For more information, see “How Firewall Filters Test a Packet's Protocol” on page 191.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> ■ parameter-problem: <code>ip-header-bad</code> (0), <code>required-option-missing</code> (1) ■ redirect: <code>redirect-for-host</code> (1), <code>redirect-for-network</code> (0), <code>redirect-for-tos-and-host</code> (3), <code>redirect-for-tos-and-net</code> (2) ■ time-exceeded: <code>ttl-eq-zero-during-reassembly</code> (1), <code>ttl-eq-zero-during-transit</code> (0) ■ unreachable: <code>communication-prohibited-by-filtering</code> (13), <code>destination-host-prohibited</code> (10), <code>destination-host-unknown</code> (7), <code>destination-network-prohibited</code> (9), <code>destination-network-unknown</code> (6), <code>fragmentation-needed</code> (4), <code>host-precedence-violation</code> (14), <code>host-unreachable</code> (1), <code>host-unreachable-for-TOS</code> (12), <code>network-unreachable</code> (0), <code>network-unreachable-for-TOS</code> (11), <code>port-unreachable</code> (3), <code>precedence-cutoff-in-effect</code> (15), <code>protocol-unreachable</code> (2), <code>source-host-isolated</code> (8), <code>source-route-failed</code> (5)
<code>icmp-type number</code>	<p>ICMP packet type field. Normally, you specify this match in conjunction with the <code>protocol</code> match statement to determine which protocol is being used on the port. For more information, see “How Firewall Filters Test a Packet's Protocol” on page 191.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): <code>echo-reply</code> (0), <code>echo-request</code> (8), <code>info-reply</code> (16), <code>info-request</code> (15), <code>mask-request</code> (17), <code>mask-reply</code> (18), <code>parameter-problem</code> (12), <code>redirect</code> (5), <code>router-advertisement</code> (9), <code>router-solicit</code> (10), <code>source-quench</code> (4), <code>time-exceeded</code> (11), <code>timestamp</code> (13), <code>timestamp-reply</code> (14), or <code>unreachable</code> (3).</p>
<code>interface interface-name</code>	Interface on which the packet was received. You can configure a match condition that matches packets based on the interface on which they were received.
<code>interface-group group-number</code>	Interface group on which the packet was received. An interface group is a set of one or more logical interfaces. For <code>group-number</code> , specify a value from 0 through 255. For information about configuration interface groups, see “Applying Firewall Filters to Interfaces” on page 205.
<code>loss-priority level</code>	<p>Match on the packet loss priority (PLP) level. Specify a single level or multiple levels: <code>low</code>, <code>medium-low</code>, <code>medium-high</code>, or <code>high</code>.</p> <p>For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>JUNOS Class of Service Configuration Guide</i>.</p>
<code>loss-priority-except level</code>	<p>Do not match on the packet loss priority level. Specify a single level or multiple levels: <code>low</code>, <code>medium-low</code>, <code>medium-high</code>, or <code>high</code>.</p> <p>For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>JUNOS Class of Service Configuration Guide</i>.</p>
<code>packet-length bytes</code>	Length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.

Table 29: Numeric Range IPv4 Firewall Filter Match Conditions (*continued*)

Match Condition	Description
<i>port number</i>	<p>TCP or UDP source or destination port field. You cannot specify both the port match and either the destination-port or source-port match conditions in the same term.</p> <p>Normally, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. For more information, see “How Firewall Filters Test a Packet’s Protocol” on page 191.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed under destination-port.</p>
<i>precedence</i> <i>ip-precedence-field</i>	<p>IP precedence field. In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): critical-ecp (0xa0), flash (0x60), flash-override (0x80), immediate (0x40), internet-control (0xc0), net-control (0xe0), priority (0x20), or routine (0x00). You can specify precedence in hexadecimal, binary, or decimal form.</p>
<i>prefix-list name</i>	<p>Match on the destination or source prefixes in the specified list name. Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p>
<i>protocol number</i>	<p>IP protocol field. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): ah (51), egp (8), esp (50), gre (47), icmp (1), igmp (2), ipip (4), ipv6 (41), ospf (89), pim (103), rsvp (46), tcp (6), or udp (17).</p>
<i>source-mac-address</i> <i>address</i>	<p>Source MAC address of a VPLS packet.</p>
<i>source-port number</i>	<p>TCP or UDP source port field. You cannot specify the port and source-port match conditions in the same term.</p> <p>Normally, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. For more information, see “How Firewall Filters Test a Packet’s Protocol” on page 191.</p> <p>In place of the numeric field, you can specify one of the text synonyms listed under destination-port.</p>
<i>source-prefix-list name</i>	<p>Match on the source prefixes in the specified list name. Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p>
<i>ttl number</i>	<p>IPv4 TTL number to match. Specify a TTL value from 1 through 127. This match condition is supported only on M320, M120, MX-series, and T-series routing platforms.</p>
<i>ttl-except number</i>	<p>IPv4 TTL number to avoid matching. Specify a TTL value from 1 through 127. This match condition is supported only on M320, M120, MX-series, and T-series routing platforms.</p>
<i>vlan-ether-type value</i>	<p>Match on the virtual local area network (VLAN) Ethernet type field of a VPLS packet.</p>
<i>vlan-ether-type-except</i> <i>value</i>	<p>Do not match on the VLAN Ethernet type field of a VPLS packet.</p>

Table 30: Numeric Range IPv6 Firewall Filter Match Conditions

Match Condition	Description
<i>address address</i>	<p>128-bit address that supports the standard syntax for IPv6 addresses. For more information, see the <i>JUNOS Routing Protocols Configuration Guide</i>.</p>

Table 30: Numeric Range IPv6 Firewall Filter Match Conditions (continued)

Match Condition	Description
destination-address address	128-bit address that is the final destination node address for the packet. The filter description syntax supports the text representations for IPv6 addresses as described in RFC 2373, <i>IP Version 6 Addressing Architecture</i> . For more information about IPv6 address syntax, see the <i>JUNOS Routing Protocols Configuration Guide</i> .
destination-port number	<p>TCP or UDP destination port field. You cannot specify both the port and destination-port match conditions in the same term.</p> <p>Normally, you specify this match in conjunction with the next-header match statement to determine which protocol is being used on the port. For more information, see “How Firewall Filters Test a Packet’s Protocol” on page 191.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), xmcp (177), zephyr-clt (2103), or zephyr-hm (2104).</p>
destination-prefix-list name	Match on the destination prefixes in the specified list name. Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.
forwarding-class <i>class</i>	Match on the forwarding class. Specify assured-forwarding , best-effort , expedited-forwarding , and network-control .
icmp-code <i>number</i>	<p>ICMP code field. This value or keyword provides more specific information than icmp-type. Because the value’s meaning depends upon the associated icmp-type, you must specify icmp-type along with icmp-code. For more information, see “How Firewall Filters Test a Packet’s Protocol” on page 191.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> ■ parameter-problem: ip-header-bad (0), required-option-missing (1) ■ redirect: redirect-for-host (1), redirect-for-network (0), redirect-for-tos-and-host (3), redirect-for-tos-and-net (2) ■ time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) ■ unreachable: communication-prohibited-by-filtering (13), destination-host-prohibited (10), destination-host-unknown (7), destination-network-prohibited (9), destination-network-unknown (6), fragmentation-needed (4), host-precedence-violation (14), host-unreachable (1), host-unreachable-for-TOS (12), network-unreachable (0), network-unreachable-for-TOS (11), port-unreachable (3), precedence-cutoff-in-effect (15), protocol-unreachable (2), source-host-isolated (8), source-route-failed (5)

Table 30: Numeric Range IPv6 Firewall Filter Match Conditions (*continued*)

Match Condition	Description
<code>icmp-type number</code>	<p>ICMP packet type field. Normally, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. For more information, see “How Firewall Filters Test a Packet’s Protocol” on page 191.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): echo-reply (0), echo-request (8), info-reply (16), info-request (15), mask-request (17), mask-reply (18), parameter-problem (12), redirect (5), router-advertisement (9), router-solicit (10), source-quench (4), time-exceeded (11), timestamp (13), timestamp-reply (14), or unreachable (3).</p>
<code>interface-group group-number</code>	Interface group on which the packet was received. An interface group is a set of one or more logical interfaces. For information about configuration interface groups, see “Applying Firewall Filters to Interfaces” on page 205.
<code>loss-priority level</code>	<p>Match on the packet loss priority (PLP) level. Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>JUNOS Class of Service Configuration Guide</i>.</p>
<code>loss-priority-except level</code>	<p>Do not match on the packet loss priority level. Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>JUNOS Class of Service Configuration Guide</i>.</p>
<code>next-header bytes</code>	<p>8-bit IP protocol field that identifies the type of header immediately following the IPv6 header. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): ah (51), dstops (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmpv6 (1), igmp (2), ipip (4), ipv6 (41), no-next-header (59), ospf (89), pim (103), routing (43), rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).</p>
<code>packet-length bytes</code>	Length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.
<code>port number</code>	<p>TCP or UDP source or destination port field. You cannot specify both the port match and either the destination-port or source-port match conditions in the same term.</p> <p>Typically, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. For more information, see “How Firewall Filters Test a Packet’s Protocol” on page 191.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed under destination-port.</p>
<code>prefix-list name</code>	Match on the source or destination prefixes in the specified list name. Specify the name of a list defined at the <code>[edit routing-options prefix-list prefix-list-name]</code> hierarchy level.
<code>source-address address</code>	Address of the source node sending the packet; 128 bits in length. The filter description syntax supports the text representations for IPv6 addresses as described in RFC 2373. For more information about IPv6 address syntax, see the <i>JUNOS Routing Protocols Configuration Guide</i> .

Table 30: Numeric Range IPv6 Firewall Filter Match Conditions (continued)

Match Condition	Description
source-port <i>number</i>	<p>TCP or UDP source port field. You cannot specify the port and source-port match conditions in the same term.</p> <p>Normally, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. For more information, see “How Firewall Filters Test a Packet’s Protocol” on page 191.</p> <p>In place of the numeric field, you can specify one of the text synonyms listed under destination-port.</p>
source-prefix-list <i>name</i>	Match on the source prefixes in the specified prefix list. Specify a prefix list name defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.
traffic-class <i>number</i>	<p>8-bit field that specifies the class-of-service (CoS) priority of the packet. The traffic-class field is used to specify a DiffServ code point (DSCP) value. The numerical value cannot be greater than 0x3f.</p> <p>This field was previously used as the ToS field in IPv4. However, the semantics of this field (for example, DSCP) are identical to IPv4.</p>
tcp-flags <i>flags</i>	<p>One or more of the following TCP flags:</p> <ul style="list-style-type: none"> ■ bit-name: fin, syn, rst, push, ack, urgent You can string multiple flags using logical operators. ■ numerical value: 0x01 through 0x20 ■ text synonym: tcp-established, tcp-initial <p>Configuring the tcp-flags match condition requires that you configure the next-header tcp match condition.</p>

Table 31: Numeric Range Layer 2 Bridging Filter Match Conditions (MX-series Only)

Match Condition	Description
destination-mac-address <i>address</i>	Destination media access control (MAC) address of a Layer 2 packet in a bridging environment.
destination-port <i>number</i>	TCP or UDP destination port field. You cannot specify both the port and destination-port match conditions in the same term.
dscp <i>number</i>	<p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant 6 bits of this byte form the DSCP. For more information, see the <i>JUNOS Class of Service Configuration Guide</i>.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p>
ether-type <i>value</i>	Match on the Ethernet type field of a Layer 2 packet in a bridging environment.
ether-type-except <i>value</i>	Do not match on the Ethernet type field of a Layer 2 packet.
forwarding class <i>class</i>	Match on the forwarding class. Specify assured-forwarding , best-effort , expedited-forwarding , or network-control .

Table 31: Numeric Range Layer 2 Bridging Filter Match Conditions (MX-series Only) (continued)

<code>forwarding-class-except value</code>	Match on the Ethernet type field of a Layer 2 packet environment. Specify <code>assured-forwarding</code> , <code>best-effort</code> , <code>expedited-forwarding</code> , or <code>network-control</code> .
<code>icmp-code number</code>	ICMP code field. The value or keyword provides more specific information than <code>icmp-type</code> . Because the value's meaning depends on the associated <code>icmp-type</code> , you must specify <code>icmp-type</code> along with <code>icmp-code</code> .
<code>icmp-type number</code>	ICMP packet type field. Normally, you specify this match in conjunction with the <code>protocol</code> match statement to determine which protocol is being used on the port.
<code>interface-group group-number</code>	Interface group on which the packet was received. An interface group is a set of one or more logical interfaces. For <code>group-number</code> , specify a value from 0 through 255.
<code>interface-group-except number</code>	Do not match on the interface group on which the packet was received.
<code>ip-address address</code>	32-bit address that supports the standard syntax for IPv4 addresses.
<code>ip-destination-address address</code>	32-bit address that is the final destination node address for the packet.
<code>ip-precedence ip-precedence-field</code>	IP precedence field. In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): <code>critical-ecp</code> (0xa0), <code>flash</code> (0x60), <code>flash-override</code> (0x80), <code>immediate</code> (0x40), <code>internet-control</code> (0xc0), <code>net-control</code> (0xe0), <code>priority</code> (0x20), or <code>routine</code> (0x00).
<code>ip-precedence-except</code>	Do not match on the IP precedence field.
<code>ip-protocol number</code>	IP protocol field.
<code>ip-source-address address</code>	IP address of the source node sending the packet.
<code>learn-vlan-1p-priority value</code>	(For bridging and VPLS protocols only) Match on the IEEE 802.1p learned VLAN priority. Specify a single value or multiple values from 0 thorough 7.
<code>learn-vlan-1p-priority-except value</code>	(For bridging and VPLS protocols only) Do not match on the IEEE 802.1p learned VLAN priority. Specify a single value or multiple values from 0 thorough 7.
<code>learn-vlan-id number</code>	Match on the VLAN identifier used for MAC learning.
<code>learn-vlan-id-except number</code>	Do not match on VLAN identifier used for MAC learning.
<code>loss-priority level</code>	Match on the packet loss priority (PLP) level. Specify a single level or multiple levels: <code>low</code> , <code>medium-low</code> , <code>medium-high</code> , or <code>high</code> . For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>JUNOS Class of Service Configuration Guide</i> .
<code>loss-priority-except level</code>	Do not match on the packet loss priority level. Specify a single level or multiple levels: <code>low</code> , <code>medium-low</code> , <code>medium-high</code> , or <code>high</code> . For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>JUNOS Class of Service Configuration Guide</i> .

Table 31: Numeric Range Layer 2 Bridging Filter Match Conditions (MX-series Only) (continued)

<i>port number</i>	TCP or UDP source or destination port. You cannot specify both the port match and either the destination-port or source-port match conditions in the same term.
<i>source-mac-address address</i>	Source MAC address of a Layer 2 packet.
<i>source-port number</i>	TCP or UDP source port field. You cannot specify the port and source-port match conditions in the same term.
<i>tcp-flags flags</i>	<p>One or more of the following TCP flags:</p> <ul style="list-style-type: none"> ■ Bit-name: fin, syn, rst, push, ack, urgent ■ Numerical value: 0x01 through 0x20 ■ Text synonym: tcp-established, tcp-initial <p>You can string together multiple flags using logical operators.</p> <p>Configuring the tcp-flags match condition requires that you configure the next-header-tcp match condition.</p>
<i>traffic-type type</i>	Match on the traffic type. Specify broadcast , multicast , unknown-unicast , or known-unicast .
<i>traffic-type-except type</i>	Do not match on the traffic type.
<i>user-vlan-1p-priority value</i>	(For bridging and VPLS protocols only) Match on the IEEE 802.1p user priority. Specify a single value or multiple values from 0 thorough 7.
<i>user-vlan-1p-priority-except value</i>	(For bridging and VPLS protocols only) Do not match on the IEEE 802.1p user priority. Specify a single value or multiple values from 0 thorough 7.
<i>user-vlan-id number</i>	Match on the first VLAN identifier that is part of the payload.
<i>user-vlan-id-except number</i>	Do not match on the first VLAN identifier that is part of the payload.
<i>vlan-ether-type value</i>	Match on the VLAN Ethernet type field of a Layer 2 bridging or VPLS packet.
<i>vlan-ether-type-except value</i>	Do not match on the VLAN Ethernet type field of a Layer 2 bridging or VPLS packet.

Specifying Address Filter Match Conditions

Address filter conditions match prefix values in a packet, such as IP source and destination prefixes. For address filter match conditions, you specify a keyword that identifies the field and one or more prefixes of that type that a packet must match. Table 32 on page 186 describes the address filter match conditions.

You can specify the address in one of the following ways:

- Single prefix. A match occurs if the value of the field matches the prefix. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
destination-address 10.0.0.0/8;
```

- Multiple prefixes. A match occurs if any one of the prefixes in the list matches the packet. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
destination-address {
    10.0.0.0/8;
    192.168.0.0/32;
}
```

The order in which you list prefixes in the list is not significant. They are all evaluated to determine whether a match occurs. If prefixes overlap, longest-match rules are used to determine whether a match occurs. Each list of prefixes contains an implicit 0/0 **except** statement, which means that any prefix that does not match any prefix in the list is explicitly considered not to match.

To specify the address prefix, use the notation *prefix/prefix-length*. If you omit *prefix-length*, it defaults to /32. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set destination-address 10
[edit firewall family family-name filter filter-name term term-name from]
user@host# show
destination-address {
    10.0.0.0/32;
}
```

To exclude a prefix, specify the string **except** after the prefix. In the following example, any addresses that fall under 192.168.10.0/8 match, except for addresses that fall under 192.168.0.0/16. All other addresses implicitly do not match this condition.

```
[edit firewall family family-name filter filter-name term term-name from]
destination-address {
    192.168.0.0/16 except;
    192.168.10.0/8;
}
```

To match all destinations except one, in this example 10.1.1.0/24, configure the match conditions as follows:

```
[edit firewall family family-name filter filter-name term term-name from]
destination-address {
    0.0.0.0/0;
    10.1.1.0/24 except;
}
```

Because the prefixes are order-independent and use longest-match rules, longer prefixes subsume shorter ones as long as they are the same type (whether you specify **except** or not). This is because anything that would match the longer prefix would also match the shorter one. In the following example:

- 172.16.1.2 matches the 172.16.0.0/10 prefix, and thus the action in the **then** statement is taken.
- 172.16.2.2 matches the 172.16.2.0/16 prefix. Because this prefix is negated (that is, marked as **except**), an explicit mismatch occurs. The next term in the filter is evaluated, if there is one. If there are no more terms, the packet is discarded.
- 10.1.2.3 does not match any of the prefixes included in the **source-address** condition. Instead, it matches the implicit 0.0.0.0/0 **except** at the end of the list, and is considered to be a mismatch.
- The 172.16.3.0/16 statement is ignored because it falls under the address 172.16.0.0/10—both are the same type.
- The 10.2.2.2 **except** statement is ignored because it is subsumed by the implicit 0.0.0.0/0 **except** statement at the end of the list.

```
[edit firewall family family-name filter filter-name term term-name from]
source-address {
  172.16.0.0/10;
  172.16.2.0/16 except;
  192.168.1.0;
  192.168.1.192/26 except;
  192.168.1.254;
  172.16.3.0/16; # ignored
  0.0.0.0/0 except; # ignored
}
```

Table 32: Address Firewall Filter Match Conditions

Match Condition	Description
<i>address prefix</i>	IP source or destination address field. You cannot specify both the address and the destination-address or source-address match conditions in the same term.
<i>destination-address prefix</i>	IP destination address field. You cannot specify the destination-address and address match conditions in the same term.
<i>destination-prefix-list prefix-list</i>	IP destination prefix list field. You cannot specify the destination-prefix-list and prefix-list match conditions in the same term.
<i>prefix-list prefix-list</i>	IP source or destination prefix list field. You cannot specify both the prefix-list and the destination-prefix-list or source-prefix-list match conditions in the same term.
<i>source-address prefix</i>	IP source address field. You cannot specify the source-address and address match conditions in the same rule.
<i>source-prefix-list prefix-list</i>	IP source prefix list field. You cannot specify the source-prefix-list and prefix-list match conditions in the same term.

You can also define a list of IP address prefixes under a *prefix-list* alias for frequent reference. You make this definition at the [edit *policy-options*] hierarchy level:


```
[edit policy-options]
policy-options {
  prefix-list prefix-list {
    address;
    address;
    address;
  }
}
```

Once you have defined a prefix list, you can use it when defining firewall filters:

```
[edit firewall family family-name filter filter-name term term-name]
from {
  source-prefix-list {
    prefix-list1;
    prefix-list2;
  }
  destination-prefix-list {
    prefix-list1;
  }
}
```

You can specify noncontiguous address prefixes in a filter term for firewall filters. Noncontiguous address prefixes are prefixes that are not adjacent or neighboring to one another. For example, in the following example, the following prefixes are noncontiguous: 0.0.0.10/0.0.0.255, 0.10.0.10/0.255.0.255, and 0.12.10.9/0.255.255.255:

```
[edit firewall family inet filter filter-name]
term term-name {
  address 0.0.0.10/0.0.0.255;
  destination-address 0.10.0.10/0.255.0.255;
  source-address 0.12.10.9/0.255.255.255 except;
}
```



NOTE: Noncontiguous address prefixes are valid only for IPv4 filters. IPv6 filters do not support noncontiguous address prefixes.

You can also specify a netmask value rather than a prefix length, for example:

```
[edit firewall family inet filter filter-name]
term term-name {
  address 10.0.0.10/255.0.0.255;
}
```

The prefix notation shown matches any address with a first and last octet of 10. The address and netmask are separated by a forward slash (/). The second and third bytes of the prefix can be any value from 0 through 255.



NOTE: When a firewall filter term includes the **from address address** match condition and a subsequent term includes the **from source-address address** match condition for the same address, packets may be processed by the latter term before they are evaluated by any intervening terms. Therefore, packets that should be rejected by the intervening terms may be accepted, or packets that should be accepted may be rejected.

To prevent this from occurring, we recommend you do the following. For every firewall filter term that contains the **from address address** match condition, replace that term with two separate terms: one that contains the **from source-address address** match condition, and another that contains the **from destination-address address** match condition.

For more information about prefixes, see the *JUNOS Routing Protocols Configuration Guide*.

Specifying Bit-Field Filter Match Conditions

Bit-field filter conditions match packet fields if particular bits in those fields are or are not set. You can match the IP options, TCP flags, and IP fragmentation fields. For bit-field filter match conditions, you specify a keyword that identifies the field and tests to determine that the option is present in the field. Table 33 on page 189 describes the bit-field match conditions.



NOTE: The JUNOS software does not automatically check the first fragment bit when matching TCP flags. To include the first fragment bit, include the **fragment-offset** match condition described in Table 29 on page 176.

To specify the bit-field value to match, enclose the value in quotation marks (“ ”). For example, a match occurs if the **RST** bit in the TCP flags field is set:

```
tcp-flags “rst”;
```

Generally, you specify the bits being tested using keywords. Bit-field match keywords always map to a single bit value. You also can specify bit fields as hexadecimal or decimal numbers.

To negate a match, precede the value with an exclamation point. For example, a match occurs only if the **RST** bit in the TCP flags field is *not* set:

```
tcp-flags “!rst”;
```

To match multiple bit-field values, use the logical operators list in Table 34 on page 190. The operators are listed in order, from highest precedence to lowest precedence. Operations are left-associative.

As an example of a logical AND operation, in the following, a match occurs if the packet is the initial packet on a TCP session:

```
tcp-flags "syn & !ack";
```

As an example of a logical OR operation, in the following, a match occurs if the packet is *not* the initial packet on a TCP session:

```
tcp-flags "!syn | ack";
```

As an example of grouping, in the following, a match occurs for any packet that is either a TCP reset or is not the initial packet in the session:

```
tcp-flags "!(syn & !ack) | rst";
```

When you specify a numeric value that has more than one bit set, the value is treated as a logical AND of the set bits. For example, the following two values are the same and a match occurs only if either bit 0x01 or 0x02 is not set:

```
tcp-flags "!0x3";
tcp-flags "!(0x01 & 0x02)";
```

You can use text synonyms to specify some common bit-field matches. You specify these matches as a single keyword. For example:

```
tcp-established;
```

Table 33: Bit-Field Firewall Filter Match Conditions

Match Condition	Description
Conditions with Variables	
<code>fragment-flags <i>number</i></code>	IP fragmentation flags. In place of the numeric field value, you can specify one of the following keywords (the field values are also listed): <code>dont-fragment</code> (0x4000), <code>more-fragments</code> (0x2000), or <code>reserved</code> (0x8000).
<code>ip-options <i>number</i></code>	IP options. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): <code>any</code> , <code>loose-source-route</code> (131), <code>record-route</code> (7), <code>router-alert</code> (148), <code>security</code> (130), <code>stream-id</code> (136), <code>strict-source-route</code> (137), or <code>timestamp</code> (68).
<code>tcp-flags <i>number</i></code>	TCP flags. Normally, you specify this match in conjunction with the <code>protocol</code> match statement to determine which protocol is being used on the port. For more details, see “How Firewall Filters Test a Packet’s Protocol” on page 191. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): <code>ack</code> (0x10), <code>fin</code> (0x01), <code>push</code> (0x08), <code>rst</code> (0x04), <code>syn</code> (0x02), or <code>urgent</code> (0x20).
Text Synonyms	
<code>first-fragment</code>	First fragment of a fragmented packet. This condition does not match unfragmented packets.
<code>is-fragment</code>	This condition matches if the packet is a trailing fragment; it does not match the first fragment of a fragmented packet. To match both first and trailing fragments, you can use two terms.
<code>tcp-established</code>	TCP packets other than the first packet of a connection. This is a synonym for " <code>(ack rst)</code> ". This condition does not implicitly check that the protocol is TCP. To check this, specify the <code>protocol tcp</code> match condition.

Table 33: Bit-Field Firewall Filter Match Conditions (*continued*)

Match Condition	Description
tcp-initial	First TCP packet of a connection. This is a synonym for "(syn & !ack)". This condition does not implicitly check that the protocol is TCP. To check this, specify the <code>protocol tcp</code> match condition.

Table 34: Bit-Field Logical Operators

Logical Operator	Description
(...)	Grouping
!	Negation
& or +	Logical AND
or ,	Logical OR

Specifying Class-Based Filter Match Conditions

Class-based filter conditions match packet fields based on source class or destination class. A source class is a set of source prefixes grouped together and given a class name. A destination class is a set of destination prefixes grouped together and given a class name.

You can specify the source class in the following way:

```
[edit firewall filter inet filter-name term term-name]
from {
  source-class class-name;
}
```

You can specify the destination class in the following way:

```
[edit firewall filter inet filter-name term term-name]
from {
  destination-class class-name;
}
```

You can specify a source class or destination class for an output firewall filter. Although you can specify a source class and destination class for an input firewall filter, the counters are incremented only if the firewall filter is applied on the output interface.

The class-based filter match condition works only for output filters because the source class usage (SCU) and destination class usage (DCU) are determined after route lookup.



NOTE: Source class usage and destination class usage are not supported on the interfaces you configure as tunnel sources. This affects only the transit packets exiting the tunnel.



NOTE: Class-based filter match conditions are not supported on the J-series Services Routers.



NOTE: Class-based filter match conditions are supported for `inet` and `inet6` address families on the M-series platforms.

Filtering Smaller Packets

By default, firewall filtering is not supported for packets that are less than 5 bytes in length. To filter packets less than 5 bytes in length, include an additional term to match the packet size.

For example, consider the following filter term:

```
term 1 {
  from {
    fragment-offset-except 0;
  }
  then {
    reject;
  }
}
```

To consider packets of 1 through 4 bytes in length, include an additional term that matches the packet size:

```
term 2 {
  from {
    packet-length [ 21 - 24 ];
  }
  then {
    reject;
  }
}
```

How Firewall Filters Test a Packet's Protocol

If you specify a port match condition or a match of the ICMP type or TCP flags field, there is no implied protocol match. If you use one of the following match conditions in a term, you should explicitly specify the protocol in the same term:

- **destination-port**—For IPv4, specify the match **protocol tcp** or **protocol udp** in the same term. For IPv6, specify the match **next-header tcp** or **next-header udp** in the same term.
- **icmp-code**—Specify the match **protocol icmp** in the same term.
- **icmp-type**—Specify the match **protocol icmp** in the same term.
- **port**—Specify the match **protocol tcp** or **protocol udp** in the same term.
- **source-port**—Specify the match **protocol tcp** or **protocol udp** in the same term.
- **tcp-flags**—Specify the match **protocol tcp** in the same term.

When examining match conditions, the policy framework software tests only the specified field itself. The software does not also test the IP header to determine that the packet is indeed an IP packet.

If you do not explicitly specify the protocol, when using the fields listed previously, design your filters carefully to ensure that they are performing the expected matches. If you specify a match of **destination-port ssh**, the policy framework software deterministically matches any packets that have a value of 22 in the 2-byte field that is 2 bytes beyond the end of the IP header, without ever checking the IP protocol field.

Example: Do Not Test Packet Protocol

The first term matches all packets except for TCP and UDP packets, so only TCP and UDP packets are evaluated by the third term (**term test-a-port**):

```
[edit]
firewall {
  family inet {
    filter test-filter {
      term all-but-tcp-and-udp {
        from {
          protocol-except [tcp udp];
        }
        then accept;
      }
      term test-an-address {
        from {
          address 192.168/16;
        }
        then reject;
      }
      term test-a-port {
        from {
          destination-port [ssh dns];
        }
        then accept;
      }
      term dump-anything-else {
        then reject;
      }
    }
  }
}
```

```
    }
}
```

Configuring a Filter Within a Filter

You can configure a filter within the term of another filter to minimize the work needed to configure terms common to numerous filters. Each firewall filter consists of one or more *terms*. You can configure one filter with the common desired terms, and apply them to other filters. To make changes to the common desired terms, you need to make term modifications only to the filter with the common terms instead of changing terms on every filter.

To configure a filter within a filter, include the `filter` statement at the `[edit firewall filter inet filter-name term term-name]` hierarchy level:

```
term term-name {
    filter filter-name;
}
```

A filter within a filter cannot reference yet another filter. For example, the following configuration is *not* valid:

```
[edit]
firewall {
  filter filter-name {
    term t1 {
      filter filter-name2 {
        term t2 {
          filter filter-name3;
        }
      }
    }
  }
}
```

You cannot configure the `from` or `then` option under the same filter term that references a filter within a filter. For example, the following configuration is *not* valid:

```
[edit]
firewall {
  filter filter-name {
    term t1 {
      filter filter-name2 {
        then {
          accept;
        }
      }
    }
  }
}
```

The maximum number of filters within a filter is limited to 256.

Example: Configuring a Filter Within A Filter

Define a filter `common-filter` and configure it into two separate filters:

```
[edit]
firewall {
  filter common-filter {
    term t1 {
      from {
        protocol udp;
        port tftp;
      }
      then {
        log;
        discard;
      }
    }
  }
  filter filter1 {
    term term1 {
      filter common-filter;
    }
  }
  filter filter2 {
    term term1 {
      filter common-filter;
    }
  }
}
```

Examples: Defining Firewall Filters

The following examples illustrate how to define firewall filters:

- Example: Blocking Telnet and SSH Access on page 195
- Example: Blocking TFTP Access on page 195
- Example: Accepting DHCP Packets with Specific Addresses on page 196
- Example: Defining a Policer for a Destination Class on page 196
- Example: Counting IP Option Packets on page 197
- Example: Accepting OSPF Packets from Certain Addresses on page 198
- Example: Matching Packets Based on Two Unrelated Criteria on page 199
- Example: Counting Both Accepted and Rejected Packets on page 199
- Example: Blocking TCP Connections to a Certain Port Except from BGP Peers on page 200
- Example: Accepting Packets with Specific IPv6 TCP Flags on page 201
- Example: Setting a Rate Limit for Incoming Layer 2 Control Packets on page 201

Example: Blocking Telnet and SSH Access

Block telnet and SSH access to all but the 192.168.1.0/24 subnet. This filter also logs any SSH or telnet traffic attempts from other subnets to the firewall log buffer:

```
[edit]
firewall {
  family inet {
    filter local-access-control {
      term terminal-access {
        from {
          address {
            192.168.1.0/24;
          }
          protocol tcp;
          port [ssh telnet];
        }
        then accept;
      }
      term terminal-access-denied {
        from {
          protocol tcp;
          port [ssh telnet];
        }
        then {
          log;
          reject;
        }
      }
      term default-term {
        then accept;
      }
    }
  }
}
```

Example: Blocking TFTP Access

Block Trivial File Transfer Protocol (TFTP) access, logging any attempts to establish TFTP connections:

```
[edit]
firewall {
  family inet {
    filter tftp-access-control {
      term one {
        from {
          protocol udp;
          port tftp;
        }
        then {
          log;
          discard;
        }
      }
    }
  }
}
```

```

    }
  }
}

```

By default, to decrease vulnerability to denial-of-service (DoS) attacks, the JUNOS software filters and discards Dynamic Host Configuration Protocol (DHCP) or Bootstrap Protocol (BOOTP) packets that have a source address of 0.0.0.0 and a destination address of 255.255.255.255. This default filter is known as a unicast RPF check. However, some vendors' equipment automatically accepts these packets. To interoperate with other vendors' equipment, you can configure a filter that checks for both these addresses and overrides the default RPF-check filter by accepting these packets.

Example: Accepting DHCP Packets with Specific Addresses

Configure a filter (`rpf-dhcp`) that accepts DHCP packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255:

```

[edit firewall family inet]
filter rpf-dhcp {
  term dhcp {
    from {
      source-address {
        0.0.0.0/32;
      }
      destination-address {
        255.255.255.255/32;
      }
    }
    then {
      accept;
    }
  }
}

```

To apply this filter to an interface, include the `rpf-check fail-filter` statement at the `[edit interface interface-name unit logical-unit-number family family-name]` hierarchy level:

```

[edit interface interface-name unit logical-unit-number family inet]
rpf-check fail-filter rpf-dhcp;

```

Example: Defining a Policer for a Destination Class

Define a policer for destination class `class1`:

```

[edit]
firewall {
  family inet {
    filter filter1 {
      policer police-class1 {
        if-exceeding {
          bandwidth-limit 25;
          burst-size-limit 1000;
        }
      }
    }
  }
}

```

```

        then {
            discard;
        }
    }
    term term1 {
        from {
            destination-class class1;
        }
        then {
            policer police-class1;
        }
    }
}
}
}

```

Example: Counting IP Option Packets

Count individual IP option packets, but do not block any traffic. Also, log packets that have loose or strict source routing:

```

[edit]
firewall {
    family inet {
        filter ip-option-filter {
            term match-strictsource {
                from {
                    ip-options strict-source-route;
                }
                then {
                    count strict-source-route;
                    log;
                    accept;
                }
            }
            term match-loose-source {
                from {
                    ip-options loose-source-route;
                }
                then {
                    count loose-source-route;
                    log;
                    accept;
                }
            }
        }
        term match-record {
            from {
                ip-options record-route;
            }
            then {
                count record-route;
                accept;
            }
        }
        term match-timestamp {

```

```

        from {
            ip-options timestamp;
        }
        then {
            count timestamp;
            accept;
        }
    }
    term match-router-alert {
        from {
            ip-options router-alert;
        }
        then {
            count router-alert;
            accept;
        }
    }
    term match-all {
        then accept;
    }
}
}
}

```

Example: Accepting OSPF Packets from Certain Addresses

Accept only Open Shortest Path First (OSPF) packets from an address in the prefix 10.108.0.0/16, discarding all other packets with an administratively-prohibited ICMP message:

```

[edit]
firewall {
    family inet {
        filter ospf-filter {
            term term1 {
                from {
                    source-address {
                        10.108.0.0/16;
                    }
                }
                protocol ospf;
            }
        }
        term default-term {
            then {
                reject administratively-prohibited; # default reject action
            }
        }
    }
}
}

```

Example: Matching Packets Based on Two Unrelated Criteria

Match packets that are either OSPF packets or packets that come from an address in the prefix 10.108/16, and send an administratively-prohibited ICMP message for all packets that do not match:

```
[edit]
firewall {
  family inet {
    filter ospf-or-131 {
      term protocol-match {
        from {
          protocol ospf;
        }
      }
      term address-match {
        from {
          source-address {
            10.108.0.0/16;
          }
        }
      }
    }
  }
}
```

Example: Counting Both Accepted and Rejected Packets

Reject all addresses except 192.168.5.0/24. In the first term, the statement 192.168.5.2/24 except causes this address to be considered a mismatch and this address is passed to the next term in the filter. The address 0.0.0.0/0 in the first term matches all other packets, and these are counted, logged, and rejected. In the second term, all packets that passed through the first term (that is, packets whose address matches 192.168.5.2/24) are counted, logged, and accepted.

```
[edit]
firewall {
  family inet {
    filter fire1 {
      term 1 {
        from {
          address {
            192.168.5.0/24 except;
            0.0.0.0/0;
          }
        }
        then {
          count reject-pref1-1;
          log;
          reject;
        }
      }
      term 2 {
        then {
          count reject-pref1-2;
        }
      }
    }
  }
}
```

```

        log;
        accept;
    }
}
}
}
}

```

Example: Blocking TCP Connections to a Certain Port Except from BGP Peers

Block all TCP connection attempts to port 179 from all requesters except the specified BGP peers:

```

[edit]
firewall {
  family inet {
    filter bgp179 {
      term 1 {
        from {
          source-address {
            0.0.0.0/0;
          }
          source-prefix-list {
            bgp179 except;
          }
          destination-port bgp;
        }
        then {
          reject;
        }
      }
      term 2 {
        then {
          accept;
        }
      }
    }
  }
}

```

Expand the prefix list **bgp179** to include all BGP group neighbors:

```

[edit policy-options]
prefix-list bgp179 {
  apply-path "protocols bgp group <*> neighbor <*>";
}

```

Apply the filter **bgp179** to interface **lo0**:

```

[edit interfaces lo0]
root@hostname# show
unit 0 {
  family inet {
    filter {

```

```

        input bgp179;
    }
    address 10.0.0.1/32;
}

```

Example: Accepting Packets with Specific IPv6 TCP Flags

Configure a filter to match on IPv6 TCP flags:

```

[edit]
firewall {
  family inet6 {
    filter tcpfilt {
      term 1 {
        from {
          next-header tcp;
          tcp-flags syn;
        }
        then {
          count tcp_syn_pkt;
          log;
          accept;
        }
      }
    }
  }
}

```

Example: Setting a Rate Limit for Incoming Layer 2 Control Packets

Configure rate limiting for incoming Layer 2 control packets. In order to meet this requirement, you must configure an input filter with the family type any and apply this filter to the interface:

```

[edit]
firewall {
  policer p1 {
    if-exceeding {
      bandwidth-limit 5m;
      burst-size-limit 10m;
    }
    then discard;
  }
  policer p2 {
    if-exceeding {
      bandwidth-limit 40m;
      burst-size-limit 100m;
    }
    then discard;
  }
  policer p3 {
    if-exceeding {
      bandwidth-limit 600m;
      burst-size-limit 1g;
    }
  }
}

```

```

    }
    then discard;
  }
  interface-set ifset {
    fe-*;
  }
  family any {
    filter L2-filter {
      term t1 {
        from {
          interface fe-0/0/0.0;
        }
        then policer p1;
      }
      term t2 {
        from {
          interface-set ifset;
        }
        then policer p2;
      }
      term t3 {
        then policer p3;
      }
    }
  }
}
[edit]
interfaces {
  fe-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.1/30;
      }
    }
  }
  fe-1/0/0 {
    unit 0 {
      family inet {
        address 10.2.2.1/30;
      }
    }
  }
  lo0 {
    unit 0 {
      family any {
        filter {
          input L2-filter;
        }
      }
    }
  }
}
}

```


Configuring Service Filters

A service filter identifies packets on which services need to be applied. The service filter also identifies the Physical Interface Card (PIC) performing the service. To configure service filters, include the **service-filter** statement at the [edit firewall family inet] hierarchy level:

```
[edit firewall family inet]
service-filter filter-name {
  term term-name {
    from {
      match-conditions;
    }
    then {
      action;
      action-modifiers;
    }
  }
}
```



NOTE: You must specify **inet** as the address family in order to configure a service filter.

Service filters are configured the same way as firewall filters. Service filters have a subset of the match conditions and actions for firewall filters.

One of the actions must be **service** or **skip**:

- Specifying the **service** action directs packets for stateful-firewall service.
- Specifying the **skip** action let packets bypass stateful-firewall service.

For more information about configuring firewall filters, see “Configuring Firewall Filters” on page 165. For more information about services and service interfaces, see the *JUNOS Services Interfaces Configuration Guide*.

Configuring Simple Filters

Simple filters are used to support Ethernet IQ2 PICs. A simple filter is a subset of a firewall filter with the following limitations:

- The **next-term** action is not supported.
- The **except** and **protocol-except** match conditions are not supported.
- Non-contiguous masks are not supported.
- Only one **source-address** and **destination-address** prefix are allowed for each filter term.

To configure simple filters, include the **simple-filter** statement at the [edit firewall family inet] hierarchy level:

```

[edit firewall family inet]
simple-filter filter-name {
  term term-name {
    from {
      match-conditions;
    }
    then {
      action;
      action-modifiers;
    }
  }
}

```

For more information about Ethernet IQ2 PICs and related features, see the *JUNOS Services Interfaces Configuration Guide* and the *JUNOS Class of Service Configuration Guide*.

Example: Configuring a Simple Filter

Configure a simple filter to support Ethernet IQ2 PICs:

```

[edit]
firewall {
  family inet {
    simple-filter sf-1 {
      term 1 {
        from {
          source-address 172.16.0.0/16;
          destination-address 20.16.0.0/16;
          source-port 1024-9071;
        }
        then {
          forwarding-class fc-be1;
          loss-priority high;
          accept;
        }
      }
      term 2 {
        from {
          source-address 173.16.0.0/16;
          destination-address 21.16.0.0/16;
        }
        then {
          forwarding-class fc-ef1;
          loss-priority low;
          accept;
        }
      }
    }
  }
}

```

Applying Firewall Filters to Interfaces

For a firewall filter to work, you must apply it to at least one interface. To do this, include the `filter` statement when configuring the logical interface at the `[edit interfaces interface-name unit logical-unit-number family family-name]` hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family]
filter {
  input filter-name;
  input-list [ filter-names ];
  output filter-name;
  output-list [ filter-names ];
}
```

In the `input` statement, list the name of one firewall filter to be evaluated when packets are received on the interface. Input filters applied to the loopback interface, `lo0`, affect only inbound traffic destined for the Routing Engine.

In the `input-list` statement, list the names of firewall filters to be evaluated when packets are received on the interface. You can specify up to 16 firewall filters for the filter input list. In the `output-list` statement, list the names of firewall filters to be evaluated when packets are transmitted from the interface. You can specify up to 16 firewall filters for the filter output list.

In the `output` statement, list the name of one firewall filter to be evaluated when packets are transmitted on the interface. Output filters applied to the loopback interface, `lo0`, affect only outbound traffic sent from the Routing Engine.

You can apply only one input and one output firewall filter to each interface. You can use the same filter one or more times.

For more information about configuring filters on interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

When you apply a filter to an interface, it is evaluated against all the data packets passing through that interface. The exception is the loopback interface, `lo0`, which is the interface to the Routing Engine and carries no data packets. If you apply a filter to the `lo0` interface, the filter affects the local packets received or transmitted by the Routing Engine.

Filters apply to all packets entering an interface, not just the packets destined for the Routing Engine. To filter packets destined for the Routing Engine, configure the `group` statement at the `[edit interfaces interface-name unit logical-unit-number family family-name filter]` hierarchy level. For more information, see “Defining Interface Groups” on page 207.

For filters applied to data packets to function, the routing platform must contain an Internet Processor II ASIC.

You can configure the following additional properties when applying filters to interfaces:

- Configuring Interface-Specific Counters on page 206
- Defining Interface Groups on page 207

Configuring Interface-Specific Counters

When you configure a firewall filter that is applied to multiple interfaces, you can name individual counters specific to each interface. These counters enable you to easily maintain statistics on the traffic transiting the different interfaces.



NOTE: Configuration of interface-specific counters also creates separate instances of any policers you have configured for the same interface. For more information about policers, see “Policer Configuration” on page 233.

To configure interface-specific counters, include the **interface-specific** statement at the `[edit firewall family family-name filter filter-name]` hierarchy level:

```
[edit firewall family filter filter-name]
interface-specific;
```



NOTE: The counter name is restricted to 24 bytes. If the renamed counter exceeds this maximum length, the policy framework software might reject it.

Example: Configuring Interface-Specific Counters

Configure an interface-specific counter:

```
[edit firewall]
family inet {
  filter test {
    interface-specific;
    term 1 {
      from {
        address {
          10.0.0.0/12;
        }
        protocol tcp;
      }
      then {
        count sample1;
        accept;
      }
    }
  }
}
```

When you apply this filter to the input interface of `at-1/1/1.0` and the output interface of `so-2/2/2.2`, the counters are named `sample1-at-1/1/1.0-i` and `sample1-so-2/2/2.2-o`. The suffixes `-i` (input) and `-o` (output) are added to the counter names automatically.

The JUNOS software does not sample packets originating from the router. If you configure a sampling filter and apply it to the output side of an interface, then only

the transit packets going through that interface are sampled. Packets that are sent from the Routing Engine to the Packet Forwarding Engine are not sampled.

Defining Interface Groups

When applying a firewall filter, you can define an interface to be part of an *interface group*. Packets received on that interface are tagged as being part of the group. You then can match these packets using the **interface-group** match statement, as described in Table 29 on page 176.

To define an interface to be part of an interface group, include the **group** statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family-name* filter] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family family-name filter]
group group-number;
input filter-name;
output filter-name;
```

In the **group** statement, specify the interface group number to be associated with the filter.

In the **input** statement, list the name of one firewall filter to be evaluated when packets are received on the interface.

In the **output** statement, list the name of one firewall filter to be evaluated when packets are transmitted on the interface.

Example: Defining Interface Groups

Create a filter that contains an interface group:

```
[edit firewall]
family inet {
  filter if-group {
    term group1 {
      from {
        interface-group 1;
        address {
          192.168.80.114/32;
        }
        protocol tcp;
        port finger;
      }
      then {
        count if-group-counter1;
        log;
        reject;
      }
    }
  }
  term group-2 {
    then {
      count if-group-counter2;
      log;
    }
  }
}
```

```

        accept;
    }
}
}

```

Assign one or more interfaces to the interface group referenced in the filter:

```

[edit interfaces]
fxp0 {
  unit 0 {
    family inet {
      filter {
        group 1;
      }
      address 192.168.5.38/24;
    }
  }
}

```

Apply the filter that contains an interface group:

```

[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      filter {
        input if-group;
        group 1;
      }
      address 10.0.0.1/32;
      address 192.168.77.1/32;
    }
  }
}

```

Configuring Firewall Filters for Logical Systems

You can configure a separate set of firewall filters for each logical system on the router. To configure a firewall filter for a logical system, you must perform at least the following tasks:

- Configure firewall filters for the logical system—To configure firewall filters for the logical system, include the **firewall** statement at the `[edit logical-systems logical-system-name]` hierarchy level:

```

[edit logical-systems logical-system-name]
firewall {
  family family-name {
    filter filter-name {
      accounting-profile name;
      interface-specific;
      term term-name {
        from {
          match-conditions;

```

```

    }
    then {
        action;
        action-modifiers;
    }
}
}
}
}

```

- Apply firewall filters to interfaces in the logical system—To have the firewall filter take effect, you must apply it to an interface in the logical system by including the `filter` statement at the `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family-name]` hierarchy level:

```

[edit logical-systems logical-system-name interfaces interface-name unit
  logical-unit-number family family-name]
filter {
    input filter-name;
    output filter-name;
}

```

To identify firewall objects configured under logical systems, operational `show` commands and firewall-related SNMP MIB objects include a `__logical-system-name/` prefix in the object name. For example, firewall objects configured under the `ls1` logical system include an `__ls1/` prefix.

This section includes the following topics:

- Guidelines for Firewall Configuration in Logical Systems on page 209
- Unsupported Configuration Statements, Actions, and Action Modifiers on page 216

Guidelines for Firewall Configuration in Logical Systems

As a general rule, firewall filters configured under a logical system must be complete and self-contained. Typically, the filters cannot reference firewall elements configured at the `[edit firewall]` hierarchy level or at another `[edit logical-systems logical-system-name]` hierarchy level. If no firewall filters are configured for a logical system, the firewall filters at the `[edit firewall]` hierarchy level are applied.

In some situations, `firewall` statements that are valid under the `[edit firewall]` hierarchy are not supported under the `[edit logical-systems logical-system-name firewall]` hierarchy. There are three scenarios to consider:

- Scenario 1. An object in the firewall hierarchy references another object in the hierarchy; for example, when a firewall filter references a firewall policer.
- Scenario 2. An object outside the firewall references an object inside the firewall hierarchy; for example, a firewall filter is applied to an interface.
- Scenario 3. An object in the firewall hierarchy references an object outside the firewall hierarchy; for example, when a firewall filter references a prefix list (defined under the `[edit policy-options]` hierarchy).

This section includes the following topics:

- Scenario 1: Firewall Objects That Reference Other Firewall Objects on page 210
- Scenario 2: Referencing Firewall Objects from Outside the Firewall Configuration on page 211
- Scenario 3: Firewall Objects That Reference Objects Outside the Firewall Configuration on page 215

Scenario 1: Firewall Objects That Reference Other Firewall Objects

If a firewall object references a subordinate object (for example, a policer or prefix list), that subordinate object must be defined within the firewall object. For example, if a firewall filter configuration references a policer, that policer must be configured under the same firewall object as the filter. This rule applies even if the same policer is configured under the main firewall configuration or if the same policer is configured as part of a firewall in another logical system.

In this example, the **filter1** filter references the **pol1** policer. Both **filter1** and **pol1** are defined under the same firewall object. This configuration is valid. If **pol1** was defined under another firewall object, the configuration would not be valid.

```
[edit]
logical systems {
  ls1 {
    firewall {
      policer pol1 {
        if-exceeding {
          bandwidth-limit 401k;
          burst-size-limit 50k;
        }
        then discard;
      }
      filter filter1 {
        term one {
          from {
            source-address 12.1.0.0/16;
          }
          then {
            reject host-unknown;
          }
        }
        term two {
          from {
            source-address 12.2.0.0/16;
          }
          then policer pol1;
        }
      }
    }
  }
}
```


Scenario 2: Referencing Firewall Objects from Outside the Firewall Configuration

When an object is configured within a logical system (but is not included in the firewall configuration for the logical system) and that object references a firewall object, the following logic is used to resolve the configuration:

- If firewall configuration statements are defined within the same logical system, the [edit logical-systems *logical-system-name* firewall] hierarchy is searched to resolve the configuration. The main [edit firewall] hierarchy is *not* searched.
- If no firewall configuration statements are defined within the same logical system, the firewall configuration defined at the [edit firewall] hierarchy level is searched to resolve the configuration. This search option is provided for legacy purposes. The main [edit firewall] hierarchy is searched only if firewall configuration statements are *not* defined within the same logical system.
- Firewall configurations that belong to other logical systems are *not* searched.

In the following example, the filter **fred** is applied to an interface in the logical system **ls1**. However, **fred** is defined in the main firewall configuration instead of in the **ls1** firewall configuration. Therefore, in this first example, the configuration is not valid.

```
[edit]
logical-systems {
  ls1 {
    interfaces {
      fe-0/3/2 {
        unit 0 {
          family inet {
            filter {
              input-list [ filter1 fred ];
            }
          }
        }
      }
    }
  }
  firewall {
    policer pol1 {
      if-exceeding {
        bandwidth-limit 401k;
        burst-size-limit 50k;
      }
      then discard;
    }
    filter filter1 {
      term one {
        from {
          source-address 12.1.0.0/16;
        }
        then {
          reject host-unknown;
        }
      }
      term two {
```

```

    from {
        source-address 12.2.0.0/16;
    }
    then policer pol1;
}
}
}
}
}
firewall {
    policer pol1 {
        if-exceeding {
            bandwidth-limit 701k;
            burst-size-limit 70k;
        }
        then discard;
    }
    family inet {
        filter fred {
            term one {
                from {
                    source-address 11.1.0.0/16;
                }
                then {
                    log;
                    reject host-unknown;
                }
            }
        }
    }
}

```

To fix this example, define filter **fred** under logical system **ls1**. In this case, the filter **fred** applied to interface **fe-0/3/2** looks for source address **10.1.0.0/16** rather than **11.1.0.0/16**.

```
[edit]
logical-systems {
  ls1 {
    interfaces {
      fe-0/3/2 {
        unit 0 {
          family inet {
            filter {
              input-list [ filter1 fred ];
            }
          }
        }
      }
    }
  }
}

firewall {
  policer pol1 {
    if-exceeding {
      bandwidth-limit 401k;
      burst-size-limit 50k;
    }
  }
}
```

```

    }
    filter filter1 {
        term one {
            from {
                source-address 12.1.0.0/16;
            }
            then {
                reject host-unknown;
            }
        }
        term two {
            from {
                source-address 12.2.0.0/16;
            }
            then policer pol1;
        }
    }
} family inet {
    filter fred {
        term one {
            from {
                source-address 10.1.0.0/16;
            }
            then {
                log;
                reject host-unknown;
            }
        }
    }
}
}
}
}
}
firewall {
    policer pol1 {
        if-exceeding {
            bandwidth-limit 701k;
            burst-size-limit 70k;
        }
        then discard;
    }
}
family inet {
    filter fred {
        term one {
            from {
                source-address {
                    11.1.0.0/16;
                }
            }
            then {
                log;
                reject host-unknown;
            }
        }
    }
}
}
```

```

    }
  }

```

If, however, the [edit logical-systems logical-system-name] hierarchy does not contain any `firewall` statements, then the main firewall configuration is used for any filter or policer references. For example, the following configuration is also allowed:

```

[edit]
logical-systems {
  ls1 {
    interfaces {
      fe-0/3/2 {
        unit 0 {
          family inet {
            filter {
              input-list [ filter1 fred ];
            }
          }
        }
      }
    }
  }
}

firewall {
  policer pol1 {
    if-exceeding {
      bandwidth-limit 701k;
      burst-size-limit 70k;
    }
    then discard;
  }
  family inet {
    filter fred {
      term one {
        from {
          source-address 11.1.0.0/16;
        }
        then {
          log;
          reject host-unknown;
        }
      }
    }
  }
  filter filter1 {
    term one {
      from {
        source-address 12.1.0.0/16;
      }
      then {
        reject host-unknown;
      }
    }
    term two {
      from {
        source-address 12.2.0.0/16;
      }
    }
  }
}

```

```

        then policer pol1;
    }
}
}

```

Scenario 3: Firewall Objects That Reference Objects Outside the Firewall Configuration

In many cases, a firewall configuration references objects outside the firewall configuration. As a general rule, the referenced object must be defined under the same logical system as the referencing object. However, there are cases when the configuration of the referenced object is not supported at the `[edit logical-systems logical-system-name]` hierarchy level.

In the following example, the service filter `inetsf1` references prefix list `prefix1`. The service set `fred` cannot be defined under the logical system `lr1`. In this case, the `[edit services]` hierarchy is searched for the definition of `fred`. This configuration is allowed because the `[edit logical-systems logical-system name]` hierarchy already had the capability to reference service sets outside the logical system hierarchy.

```

[edit]
logical-systems {
  ls1 {
    interfaces {
      fe-0/3/2 {
        unit 0 {
          family inet {
            service {
              input {
                service-set fred service-filter lr1inetsf1;
              }
            }
          }
        }
      }
    }
  }
  policy-options {
    prefix-list prefix1 {
      1.1.0.0/16;
      1.2.0.0/16;
      1.3.0.0/16;
    }
  }
  firewall {
    policer pol1 {
      if-exceeding {
        bandwidth-limit 401k;
        burst-size-limit 50k;
      }
      then discard;
    }
    filter filter1 {
      term one {
        from {

```

```

    source-address 12.1.0.0/16;
  }
  then {
    reject host-unknown;
  }
}
term two {
  from {
    source-address 12.2.0.0/16;
  }
  then policer pol1;
}
}
family inet {
  service-filter inetsf1 {
    term term1 {
      from {
        source-prefix-list {
          prefix1;
        }
      }
      then count prefix1;
    }
  }
}
}
}
}
services {
  service-set fred {
    max-flows 100;
    interface-service {
      service-interface sp-1/2/0.0;
    }
  }
}
}
```

Unsupported Configuration Statements, Actions, and Action Modifiers

Table 35 on page 217 includes a list of additional statements that are supported at the [edit firewall] hierarchy level but not supported at the [edit logical-systems *logical-system-name* firewall] hierarchy level.

Table 35: Unsupported Firewall Statements for Logical Systems

Statement	Example	Description
accounting-profile	<pre>[edit] logical-systems { ls1 { firewall { family inet { filter myfilter { accounting-profile fw_profile; ... term accept-all { then { count counter1; accept; } } } } } } }</pre>	<p>In this example, the <code>accounting-profile</code> statement is not allowed because the accounting profile <code>fw-profile</code> is configured under the <code>[edit accounting-options]</code> hierarchy.</p>
load-balance-group	<pre>[edit] logical-systems { ls1 { firewall { load-balance-group lb-group-name { next-hop-group [nh-group-names]; } } } }</pre>	<p>This configuration is not allowed because the <code>next-hop-group</code> <i>nh-group-names</i> statement must be configured at the <code>[edit forwarding-options next-hop-group]</code> hierarchy level—outside the <code>[edit logical-systems logical-system-name firewall]</code> hierarchy.</p> <p>Currently, the <code>forwarding-options dhcp-relay</code> statement is the only forwarding option supported for logical systems.</p>
virtual-channel	<pre>[edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { virtual-channel sammy; } } } } } } }</pre>	<p>This configuration is not allowed because the virtual channel <code>sammy</code> refers to an object defined at the <code>[edit class-of-service]</code> hierarchy level and class of service is not supported for logical systems.</p>

Table 36 on page 218 includes a list of the firewall filter actions and action modifiers that are supported at the [edit firewall] hierarchy level, but not supported at the [edit logical-systems *logical-system-name* firewall] hierarchy level.

Table 36: Unsupported Firewall Actions and Action Modifiers for Logical Systems

Acton or Action Modifier	Example	Description
analyzer	<pre>[edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { analyzer; } } } } } } }</pre>	(EX-series swtiches) Because the analyzer action relies on a configuration defined at the [edit ethernet-switching-options] hierarchy level, this action is not supported.
ipsec-sa	<pre>[edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { ipsec-sa barney; } } } } } } }</pre>	Because the ipsec-sa action modifier references barney , a security association defined outside the local logical system, this action is not supported.

Table 36: Unsupported Firewall Actions and Action Modifiers for Logical Systems *(continued)*

Action or Action Modifier	Example	Description
logical-system	<pre>[edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { logical-system fred; } } } } } } }</pre>	Because the <code>logical-system</code> action refers to <code>fred</code> , a logical system defined outside the local logical system, this action is not supported.
next-hop-group	<pre>[edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { next-hop-group fred; } } } } } } }</pre>	Because the <code>next-hop-group</code> action refers to <code>fred</code> , an object defined at the <code>[edit forwarding-options next-hop-group]</code> hierarchy level, this action is not supported.

Table 36: Unsupported Firewall Actions and Action Modifiers for Logical Systems *(continued)*

Action or Action Modifier	Example	Description
port-mirror	<pre> [edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { port-mirror; } } } } } } } </pre>	Because the port-mirror action relies on a configuration defined at the [edit forwarding-options port-mirroring] hierarchy level, this action is not supported.
sample	<pre> [edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { sample; } } } } } } } </pre>	In this example, the sample action depends on the sampling configuration defined under the [edit forwarding-options] hierarchy. Therefore, the sample action is not supported.

Table 36: Unsupported Firewall Actions and Action Modifiers for Logical Systems (continued)

Action or Action Modifier	Example	Description
syslog	<pre> [edit] logical-systems { ls1 { firewall { family inet { filter icmp-syslog { term icmp-match { from { address { 192.168.207.222/32; } protocol icmp; } then { count packets; syslog; accept; } } term default { then accept; } } } } } } </pre>	<p>In this example, there must be at least one system log (<i>system syslog file filename</i>) with the firewall facility enabled for the icmp-syslog filter's logs to be stored.</p> <p>Because this firewall configuration relies on a configuration outside the logical system, the <i>syslog</i> action modifier is not supported.</p>

Configuring Accounting

Juniper Networks routing platforms can collect various kinds of data about traffic passing through the routing platform. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- Fields used in the accounting records
- Number of files that the routing platform retains before discarding, and the number of bytes per file
- Polling period that the system uses to record the data

To implement an accounting profile, you must configure the profile and then apply it to an interface or firewall filter. For more information, see the *JUNOS Network Management Configuration Guide*.

Configuring a Firewall Filter Accounting Profile

There are several types of accounting profiles: interface, firewall filter, destination class, and Routing Engine. To configure an accounting profile, include statements at the [edit accounting-options] hierarchy level. To activate firewall filter profiles, you must reference them at the [edit firewall family *family-name*] hierarchy level. If you

reference the same profile name from both a firewall filter and an interface in the same configuration, it causes an error.

The following example shows accounting profile **fw_profile** for the firewall filter **myfilter**. For more information about configuring accounting profiles, see the *JUNOS Network Management Configuration Guide*.

```
[edit]
accounting-options {
  filter-profile fw_profile {
    file fw_accounting;
    interval 60;
    counters {
      counter1;
      counter2;
      counter3;
    }
  }
}
```

To apply the **fw_profile** accounting profile to a firewall filter, include the **accounting-profile** statement as shown:

```
[edit]
firewall {
  family inet {
    filter myfilter {
      accounting-profile fw_profile;
      ...
      term accept-all {
        then {
          count counter1;
          accept;
        }
      }
    }
  }
}
```

Configuring Filter-Based Forwarding

You can configure filters to classify packets based on source address and specify the forwarding path the packets take within the router by configuring a filter on the ingress interface. For example, you can use this filter for applications to differentiate traffic from two clients that have a common access layer (for example, a Layer 2 switch) but are connected to different Internet service providers (ISPs). When the filter is applied, the router can differentiate the two traffic streams and direct each to the appropriate network. Depending on the media type the client is using, the filter can use the source IP address to forward the traffic to the corresponding network through a tunnel. You can also configure filters to classify packets based on IP protocol type or IP precedence bits.



NOTE: Source-class usage filter matching and unicast reverse-path forwarding checks are not supported on an interface configured with filter-based forwarding (FBF).

You can also forward packets based on output filters by configuring a filter on the egress interfaces. In the case of port mirroring, it is useful for port-mirrored packets to be distributed to multiple monitoring PICs and collection PICs based on patterns in packet headers. FBF on the port-mirroring egress interface must be configured.

Packets forwarded to the output filter have been through at least one route lookup when an FBF filter is configured on the egress interface. After the packet is classified at the egress interface by the FBF filter, it is redirected to another routing table for further route lookup.

Filter-based forwarding is supported for IPv4 and IPv6.

To direct traffic meeting defined match conditions to a specific routing instance, include the `routing-instance` filter action:

```
routing-instance routing-instance;
```

For IPv4 traffic, include the action at the [edit firewall family inet filter *filter-name* term *term-name* then] hierarchy level. For IPv6 traffic, include the action at the [edit firewall family inet6 filter *filter-name* term *term-name* then] hierarchy level. For MPLS traffic, configure the filter terms at the [edit firewall family mpls filter *filter-name* term *term-name* then] hierarchy level.

The `routing-instance` filter action accepts the traffic meeting the match conditions and directs it to the routing instance named in *routing-instance*. For information about forwarding instances and routing instances, see the *JUNOS Routing Protocols Configuration Guide*.

To complete the configuration, you must also create a routing table group that adds interface routes to the following routing instances:

- Routing instance named in the action
- Default routing table inet.0

You create a routing table group to resolve the routes installed in the routing instance to directly connected next hops on that interface. For more information on routing table groups and interface routes, see the *JUNOS Routing Protocols Configuration Guide*.

Examples: Configuring Filter-Based Forwarding

Configure a filter to direct traffic to ISP1 or ISP2 based on source address matching:

```
[edit firewall]
family inet {
  filter classify-customers {
    term isp1-customers {
      from {
```

```

        source-address 10.1.1.0/24;
        source-address 10.1.2.0/24;
    }
    then {
        routing-instance isp1-route-table;
    }
}
term isp2-customers {
    from {
        source-address 10.2.1.0/24;
        source-address 10.2.2.0/24;
    }
    then {
        routing-instance isp2-route-table;
    }
}
term default {
    then {
        accept;
    }
}
}
}

```

Configure a filter-based forwarding (FBF) filter for family `inet6`:

```

[edit]
firewall {
    family inet6 {
        filter ftf_fbf {
            term 0 {
                from {
                    source-address {
                        ::10.34.1.0/120;
                    }
                }
                then {
                    count ce1;
                    log;
                    routing-instance ce1;
                }
            }
            term 1 {
                from {
                    source-address {
                        ::10.34.2.0/120;
                    }
                }
                then {
                    count ce2;
                    log;
                    routing-instance ce2;
                }
            }
            term default {
                then {

```

```

count default;
accept;
}
}
}
}
}
}

```

Configuring Forwarding Table Filters

The following sections describe the following topics:

- Overview of Forwarding Table Filters on page 225
- Configuring a Forwarding Table Filter on page 225

Overview of Forwarding Table Filters

You set up a forwarding table filter in essentially the same way a firewall filter: you define it, then you apply it. However, you apply the filters differently:

- Instead of applying a forwarding table filter to an interface, you apply it to a forwarding table, which is associated with a routing instance and a virtual private network (VPN).
- Instead of applying input and output firewall filters, by default, you can apply an input forwarding table filter only.

All packets are subjected to the input forwarding table filter that applies to the forwarding table. A forwarding table filter controls which packets the router accepts and performs a forwarding table lookup for, thereby controlling which packets the router forwards on the interfaces.

When the router receives a packet, it determines where to forward the packet by looking in a forwarding table, which is associated with the VPN on which the packet will be sent, for the best route to the destination. The router then forwards the packet toward its destination through the appropriate interface.



NOTE: Forwarding table filtering is not supported on the interfaces you configure as tunnel sources. This affects only the transit packets exiting the tunnel.

Configuring a Forwarding Table Filter

A forwarding table filter allows you to filter data packets based on their components and to perform an action on packets that match the filter.

To configure a forwarding table filter, do the following:

1. Define a forwarding table filter:
 - a. Configure the family address type: IPv4 (*inet*), IPv6 (*inet6*), Layer 2 traffic *bridge*, or MPLS (*mpls*).
 - b. Define one or more *terms*, which are named structures in which match conditions and actions are defined.
 - c. Define a *match condition*, which is the criterion against which a bearer packet is compared; for example, the IP address of a source device or a destination device. You can specify multiple criteria in a match condition.
 - d. Define an *action*, which is what happens if all criteria match; for example, the gateway GPRS support node (GGSN) accepting the bearer packet, performing a lookup in the forwarding table, and forwarding the packet to its destination; discarding the packet; and discarding the packet and returning a rejection message. In addition to an action, you can define one or more *action modifiers*, which are actions that are taken in addition to the GGSN accepting or discarding a packet when all criteria match; for example, counting the packets and logging a packet.

For more information about configuring firewall filters, see “Configuring Firewall Filters” on page 165.

2. Apply the forwarding table filter as an input filter to a forwarding table. The forwarding table filter controls which bearer packets the router accepts and forwards.

For more information about configuring firewall filters, see “Configuring Firewall Filters” on page 165.

To define a forwarding table filter, include the **firewall** statement at the **[edit]** hierarchy level:

```
[edit]
firewall {
  family family-name {
    filter filter-name {
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
}
```

To create a forwarding table, include the **instance-type** statement at the **[edit routing-instance *instance-name*]** hierarchy level:

```
[edit]
```



```

routing-instance instance-name {
    instance-type forwarding;
}

```

To apply a forwarding table filter to a VPN routing and forwarding (VRF) table, include the `filter input` statement at the `[edit routing-instance instance-name forwarding-options family family-name]` hierarchy level:

```

[edit]
routing-instance instance-name {
    instance-type forwarding;
    forwarding-options {
        family family-name {
            filter {
                input filter-name;
            }
        }
    }
}

```

To apply a forwarding table filter to a forwarding table, include the `filter input` statement at the `[edit forwarding-options family family-name]` hierarchy level:

```

[edit forwarding-options family family-name]
filter {
    input filter-name;
}

```

To apply a forwarding table filter to the default forwarding table `inet.0`, which is not associated with a specific routing instance, include the `filter input` statement at the `[edit forwarding-options family inet]` hierarchy level:

```

[edit]
filter {
    input filter-name;
}

```

For information about the `routing-instance` and `routing-options` statements, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring Firewall Filter System Logging

System logging can be configured for the firewall filter process. You can set system logging to record messages of a particular level or all levels. The messages are sent to a system logging file.

The following is a sample system logging configuration for the firewall filter `icmp-syslog`. For more information about configuring system logging, see the *JUNOS System Basics Configuration Guide*.

```

[edit]
system {
    syslog {
        file filter {

```

```

        firewall any;
        archive no-world-readable;
    }
}

```

This causes the **syslog** to write any messages with the **syslog** facility of **firewall** to the file **/var/log/filter**. This keeps the messages out of the main system log file and makes them easier to find.

Example: Configuring Firewall Filter System Logging

Create a filter that logs and counts ICMP packets that have **192.168.207.222** as either their source or destination:

```

[edit]
firewall {
  family inet {
    filter icmp-syslog {
      term icmp-match {
        from {
          address {
            192.168.207.222/32;
          }
          protocol icmp;
        }
        then {
          count packets;
          syslog;
          accept;
        }
      }
      term default {
        then accept;
      }
    }
  }
}

```

Enter the **show log filter** command to display the results:

```

root@hostname> show log filter
Mar 20 08:03:11 hostname feb FW: so-0/1/0.0  A icmp 192.168.207.222
192.168.207.223      0      0 (1 packets)

```

This output file contains the following fields:

- **Date and Time**—Date and time at which the packet was received (not shown in the default).
- **Filter action:**
 - **A**—Accept (or next term)
 - **D**—Discard

- R—Reject
- Protocol—Packet's protocol name or number.
- Source address—Source IP address in the packet.
- Destination address—Destination IP address in the packet.



NOTE: If the protocol is ICMP, the ICMP type and code are displayed. For all other protocols, the source and destination ports are displayed.

The last two fields (both zero) are the source and destination TCP/UDP ports, respectively, and are shown for TCP or UDP packets only. This log message indicates that only one packet for this match has been detected in about a one-second interval. If packets arrive faster, the system log function compresses the information so that less output is generated, and displays an output similar to the following:

```
root@hostname> show log filter
Mar 20 08:08:45 hostname feb FW: so-0/1/0.0   A icmp 192.168.207.222
192.168.207.223      0      0 (515 packets)
```


Chapter 10

Policer Overview

Policing, or rate limiting, enables you to limit the amount of traffic that passes into or out of an interface. It is an essential component of firewall filters that is designed to thwart denial-of-service (DoS) attacks. Policing applies two types of rate limits on the traffic:

- Bandwidth—The number of bits per second permitted, on average.
- Maximum burst size—The maximum size permitted for bursts of data that exceed the given bandwidth limit.

Policing uses the *token-bucket algorithm*, which enforces a limit on average bandwidth while allowing bursts up to a specified maximum value. It offers more flexibility than the *leaky bucket algorithm* (see the *JUNOS Class of Service Configuration Guide*) in allowing a certain amount of bursty traffic before it starts discarding packets.

You can define specific classes of traffic on an interface and apply a set of rate limits to each. You can use a policer in one of two ways: as part of a filter configuration or as part of a logical interface (where the policer is applied to all traffic on that interface).

After you have defined and named a policer, it is stored as a template. You can later use the same policer name to provide the same policer configuration each time you wish to use it. This eliminates the need to define the same policer values more than once.

Chapter 11

Policer Configuration

To configure policers, you include statements at the `[edit firewall]` hierarchy level of the configuration:

```
[edit firewall]
policer policer-name {
  filter-specific;
  if-exceeding {
    bandwidth-limit bps;
    bandwidth-percent number;
    burst-size-limit bytes;
  }
  logical-bandwidth-policer;
  logical-interface-policer;
  then {
    policer-action;
  }
}
interface-set interface-set-name {
  [ interface-names ];
}
family family-name {
  filter filter-name {
    accounting-profile name;
    interface-specific;
  }
  prefix-action name {
    count;
    destination-prefix-length prefix-length;
    policer policer-name;
    source-prefix-length prefix-length;
    subnet-prefix-length prefix-length;
  }
}
load-balance-group group-name {
  next-hop-group [ group-names ];
}
three-color-policer (Configuring) name {
  action {
    loss-priority high then discard;
  }
  logical-interface-policer;
  single-rate {
    (color-aware | color-blind);
  }
}
```

```

        committed-information-rate bps;
        committed-burst-size bytes;
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-information-rate bps;
        committed-burst-size bytes;
        peak-information-rate bps;
        peak-burst-size bytes;
    }
}

```

The following sections describe the tasks required for configuring policers and provide configuration examples:

- Minimum Policer Configuration on page 234
- Configuring Policers on page 235
- Configuring Multifield Classification and Policing on page 237
- Configuring an Interface Set on page 245
- Applying an Interface Policer on page 245
- Configuring an Aggregate Policer on page 246
- Configuring a Bandwidth Policer on page 247
- Configuring a Load-Balance Group on page 248
- Configuring Tricolor Marking on page 248
- Examples: Configuring Policing on page 252

Minimum Policer Configuration

To configure a policer, you must perform at least the following tasks:

- Configure policers—To configure policers, include the **policer** statement at the [edit firewall] hierarchy level. After policers are defined, you reference them in the then clause of a term:

```

[edit firewall]
policer policer-name {
    if-exceeding {
        bandwidth-limit bps;
        bandwidth-percent number;
        burst-size-limit bytes;
    }
    then {
        policer-action;
    }
}
family family-name {
    filter filter-name {
    }
}

```


- Add actions, such as **accept**, **discard**, or **next term**, or action modifiers, such as **count** or **log**.
- Apply the policers to an interface to activate them.

The policer is applied to the packet first, and if the packet exceeds the defined limits, the actions of the **then** clause of the policer are applied. If the result of the policing action is not a discard, the remaining components of the **then** clause of the term are applied.



NOTE: If an input filter is configured on the same logical interface as the policer, the policer is executed first.

To display statistics about a filter statement policer configuration, use the **show policers** command.

Configuring Policers

You can configure a new policer for each filter or term that requires policing. To configure term-specific policers, include the **policer** statement at the **[edit firewall]** hierarchy level:

```
[edit firewall]
policer policer-name {
  if-exceeding {
    bandwidth-limit;
    bandwidth-percent number;
    burst-size-limit;
  }
  then {
  }
}
```

The following sections describe the components of the **policer** statement and provide policer configuration examples:

- Configuring Rate Limiting on page 235
- Configuring a Policer Action on page 236

Configuring Rate Limiting

To specify the rate limiting part of a policer, include an **if-exceeding** statement at the **[edit firewall policer *policer-name*]** hierarchy level:

```
[edit firewall policer]
if-exceeding {
  bandwidth-limit bps;
  bandwidth-percent number;
  burst-size-limit;
}
```

You specify the bandwidth limit in bits per second (bps). You can specify the value as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). There is no absolute minimum value for bandwidth limit, but any value below 61,040 bps will result in an effective rate of 30,520 bps. The maximum bandwidth limit is 32 gigabits per second (Gbps).

You can rate-limit traffic based upon port speed. This port speed can be specified by a bandwidth percentage in a policer. You must specify the percentage as a complete decimal number between 1 and 100.



NOTE: You cannot rate-limit based on bandwidth percentage for aggregate, tunnel, and software interfaces. The bandwidth percentage policer cannot be used for forwarding table filters. This can only be used for interface-specific filters.

The maximum burst size controls the amount of traffic bursting allowed. To determine the value for the burst-size limit, the preferred method is to multiply the bandwidth (bytes per second) of the interface on which you are applying the filter by the amount of time you allow a burst of traffic at that bandwidth to occur; for example, 5 milliseconds.

$$\text{burst size} = \text{bandwidth} \times \text{allowable time for burst traffic} / 8$$

If you do not know the interface bandwidth, you can multiply the maximum transmission unit (MTU) of the traffic on the interface by 10 to obtain a value. For example, the burst size for an MTU of 4700 would be 47,000 bytes. At minimum, burst size should be at least 10 interface MTUs. The maximum value for the burst-size limit is 100 megabits per second (Mbps).

For a sample filter configuration for rate limiting, see “Examples: Configuring Policing” on page 252.

Configuring a Policer Action

If a packet does not exceed its rate limits, it is processed further without being affected. If the packet exceeds its limits, it is handled in one of two ways, depending on what you specify:

- Discarded
- Marked for subsequent processing based on its loss priority and forwarding class

To configure a policer action, include the **then** statement at the [edit firewall policer *policer-name*] hierarchy level:

```
[edit firewall policer policer-name]  
then {  
}
```

Policer actions include one or more of the following:

- **discard**—Discard a packet that exceeds the rate limits.

- **loss-priority *level***—Set the loss priority level to **low**, **medium-low**, **medium-high**, or **high**.
- **forwarding-class *class-name***—Specify the forwarding class to any class name already configured for the forwarding class.

Example: Configuring a Policer Action

Discard any packet that exceeds a bandwidth of 300 kilobits per second (Kbps) and a burst-size limit of 500 kilobytes (KB):

```
[edit firewall]
policer p1 {
  if-exceeding {
    bandwidth-limit 300k;
    burst-size-limit 500k;
  }
  then {
    discard;
  }
}
```

Configuring Multifield Classification and Policing

Multifield classifiers take action on incoming or outgoing packets, depending whether the firewall rule is applied as an input filter or an output filter. When TCM is enabled, T-series and M320 platforms support four multifield classifier packet loss priority (PLP) designations: **low**, **medium-low**, **medium-high**, and **high**.

To configure the PLP for a multifield classifier, include the **loss-priority** statement in a policer or firewall filter that you configure at the **[edit firewall]** hierarchy level:

```
[edit firewall]
family family-name {
  filter filter-name {
    term term-name {
      from {
        match-conditions;
      }
      then {
        loss-priority (low | medium-low | medium-high | high);
        forwarding-class class-name;
      }
    }
  }
}
```

The inputs (match conditions) for a multifield classifier are one or more of the six packet header fields: destination address, source address, IP protocol, source port, destination port, or DSCP. The outputs for a multifield classifier are the forwarding class, the PLP, or both. In other words, a multifield classifier sets the forwarding class and the PLP for each packet entering or exiting the interface with a specific destination address, source address, IP protocol, source port, destination port, or DSCP.

For example, in the following configuration, the forwarding class **expedited-forwarding** and PLP **medium-high** are assigned to all IPv4 packets with the 10.1.1.0/24 or 10.1.2.0/24 source address:

```

firewall {
  family inet {
    filter classify-customers {
      term isp1-customers {
        from {
          source-address 10.1.1.0/24;
          source-address 10.1.2.0/24;
        }
        then {
          loss-priority medium-high;
          forwarding-class expedited-forwarding;
        }
      }
    }
  }
}

```

To use this classifier, you must configure the settings for the **expedited-forwarding** forwarding class at the [edit class-of-service forwarding-classes queue *queue-number* **expedited-forwarding**] hierarchy level.



NOTE: Because the policer is executed before the filter, if an input policer is also configured on the logical interface, it cannot use the forwarding class and PLP of a multifield classifier associated with the interface.

You can configure *multifield classifiers* within a firewall filter to set the packet's forwarding class and packet loss priority. You can also apply policers to packets matching some classification term. The policing action might affect the resulting forwarding class, packet loss priority, and accept or drop status. For more information, see the *JUNOS Class of Service Configuration Guide*.

To configure the forwarding class and loss priority, include the **then** statement:

```

then {
  loss-priority;
  forwarding-class class-name;
}

```

You can include the statement at the following hierarchy levels:

- [edit firewall filter *filter-name* term *term-name*]
- [edit firewall policer *policer-name*]

You can specify one or both of the following actions:

- **loss-priority**—Set the loss priority level to **low** or **high**.
- **forwarding-class**—Specify the forwarding class to any class name already configured for the forwarding class.

For more information about forwarding class and loss priority, see the *JUNOS Class of Service Configuration Guide*. For more information about policers, see the following sections:

- Configuring Filter-Specific Policers on page 239
- Configuring Prefix-Specific Actions on page 239
- Examples: Classifying Traffic on page 244

Configuring Filter-Specific Policers

You can configure filter-specific policers within the firewall configuration. Filter-specific policers allow you to configure policers and counters for a specific filter name.

When you configure the **filter-specific** statement, a single policer set is created for the entire filter. All traffic matching the terms of the firewall filter with the action **policer** goes through that single policer. The default is a term-specific policer in which a single policer set is created for each term within the filter. All traffic matching the terms of the firewall filter with the action **policer** goes through the part of the policer that is specific to that term.

To configure filter-specific policers, include the **filter-specific** statement at the [edit firewall policer *policer-name*] hierarchy level:

```
[edit firewall policer policer-name]
filter-specific;
```

If the **filter-specific** statement is not configured, then the policer defaults to a term-specific policer.

You can apply the filter-specific policers to the family **inet**.

Configuring Prefix-Specific Actions

You can configure prefix-specific actions within the firewall configuration. Prefix-specific actions allow you to configure policers and counters for specific addresses or ranges of addresses. This allows you to essentially create policers and counters on a per-prefix level.

To configure prefix-specific actions, include the **prefix-action name** statement at the [edit firewall family **inet**] hierarchy level:

```
[edit firewall family inet]
prefix-action name {
  count;
  destination-prefix-length prefix-length;
  policer policer-name;
  source-prefix-length prefix-length;
  subnet-prefix-length prefix-length;
}
```

The following formula determines the number of prefix-specific actions created:

$$\text{Number} = 2 ^ {(\text{source/destination-prefix-length} - \text{subnet-prefix-length})}$$

The **subnet-prefix-length** statement allows for more control for the flexibility offered by prefix-specific actions, allowing the policers to be more applicable and powerful. For example, if you want to filter all Transmission Control Protocol (TCP) packets and define two policers, all packets ending with 0 in the last address bit increment the first policer, while all packets ending with 1 in the address bit increment the second policer. As another example, if you want to filter all TCP packets and define 256 policers, matching is based on the last octet of the destination address field. You achieve both cases by specifying an appropriate subnet prefix length.

Prefix-specific action is supported for the IP version 4 (IPv4) **inet** address family.

To configure prefix-specific actions, include the **prefix-action** statement and specify an action name.

To enable a prefix-specific counter, include the **count** statement.

To configure the destination address range specified for a prefix-specific policer or counter, include the **destination-prefix-length** statement.

To enable a set of prefix-specific policers, include the **policer** statement and specify the **policer** name.

To configure the source address range specified for a prefix-specific policer or counter, include the **source-prefix-length** statement.

To configure the total address range of the subnet supported, include the **subnet-prefix-length** statement. The source or destination prefix length must be larger than the subnet prefix length.

Prefix-specific action applies to a specific prefix length, and not to a specific interface. You can add an interface policer polices at the aggregate level for a specific interface. You could also use the **next term** action to configure all Hypertext Transfer Protocol (HTTP) traffic to each host to transmit at 500 Kbps and have the total HTTP traffic limited to 1 Mbps.

The maximum number of policers you can configure for one subnet is 65,536. If you configure more than 65,536 policers, you receive an error message.



NOTE: J-series Services Routers do not support prefix-specific actions.

Examples: Configuring Prefix-Specific Actions

Create a prefix-specific policer operating on the source address and apply it to the input interface:

```
[edit]
firewall {
  policer host-policer {
    filter-specific;
    if-exceeding {
      bandwidth-limit bps;
```

```

        burst-size-limit bytes;
    }
    then {
        discard;
    }
}
family inet {
    prefix-action ftp-policer-set {
        count;
        destination-prefix-length 32;
        policer host-policer;
        subnet-prefix-length 24;
    }
    filter filter-ftp {
        term term1 {
            from {
                destination-address 10.10.10/24;
                destination-port ftp;
            }
            then {
                prefix-action ftp-policer-set;
            }
        }
    }
}
}

```

Filter all packets going to the /24 subnet, letting them pass to the prefix-specific action policers. In the policer set, the last octet of the source address field of the packet is used to index into the respective prefix-specific action policers.

```

[edit]
firewall {
    policer 1Mbps-policer {
        if-exceeding {
            bandwidth-limit 1m;
            burst-size-limit 63k;
        }
    }
    family inet {
        prefix-action per-source-policer {
            policer 1Mbps-policer;
            subnet-prefix-length 24;
            source-prefix-length 32;
        }
    }
    filter limit-all-hosts {
        term one {
            from {
                source-address {
                    10.10.10.0/24;
                }
            }
            then prefix-action per-source-policer;
        }
    }
}

```

```
}
```

In the preceding case, all packets are subjected to the prefix-specific action policing. The last octet of the source address field of the packet is used to index into the corresponding policer. In other words, all packets ending with 0x(xxxx0000) match the first policer and all packets ending in 0x(xxxx0001) match the second policer.

Therefore, 256 policers are created and shared by all addresses. In this case, 10.1.1.1, 10.2.2.1, 10.4.5.1 ... 10.x.x.1 share the same 1-Mbps policer; 10.1.1.2, 10.2.2.2, 10.4.5.2 ... 10.x.x.2 share another 1-Mbps policer, and so on.

Subject packets belonging to the 10.10.10.0/24 subnet are subject to policing by the prefix-specific action policers. Because 128 policers defined in the policer set, the /24 subnet can be thought of as being split into two /25 subnets, both of them sharing the same prefix-specific action set. Therefore, 10.10.10.1 and 10.10.10.129 share the same 1-Mbps policer, 10.10.10.2 and 10.10.10.130 share another 1-Mbps policer, and so on.

```
[edit]
firewall {
  policer 1Mbps-policer {
    if-exceeding {
      bandwidth-limit 1m;
      burst-size-limit 63k;
    }
  }
  family inet {
    prefix-action per-source-policer {
      policer 1Mbps-policer;
      subnet-prefix-length 25;
      source-prefix-length 32;
    }
  }
  filter limit-all-hosts {
    term one {
      from {
        source-address {
          10.10.10.0/24;
        }
      }
      then prefix-action per-source-policer;
    }
  }
}
```

Define 256 policers based on the last octet of the source address field. However, you are only allowing a subset of that to pass through the match condition. As a result, only the lower half of the set is used.

```
[edit]
firewall {
  policer 1Mbps-policer {
    if-exceeding {
```



```

        bandwidth-limit 1m;
        burst-size-limit 63k;
    }
}
family inet {
    prefix-action per-source-policer {
        policer 1Mbps-policer;
        subnet-prefix-length 24;
        source-prefix-length 32;
    }
}
filter limit-all-hosts {
    term one {
        from {
            source-address {
                10.10.10.0/25;
            }
        }
        then prefix-action per-source-policer;
    }
}
}

```

Accept packets from **10.10.10/24** and **10.11/16** subnets and subject them to policing by the same set of prefix-specific action policers. The policers are shared by packets across both subnets. There is a one-to-one correspondence between the **10.10.10/24** subnet. For **10.11/16**, there is a many-to-one correspondence, as explained in the previous examples. Each of the **10.11.0/24**, **10.11.1/24**, **10.11.2/24** ... **10.11.255/24** subnets share the same prefix-specific action set.

Thus, **10.10.10.1**, **10.11.1.1**, **10.11.2.1** ... **10.11.x.1** share the same 1-Mbps policer; **10.10.10.2**, **10.11.1.2**, **10.11.2.2** ... **10.11.x.2** share another 1-Mbps policer, and so on.

```

[edit]
firewall {
    policer 1Mbps-policer {
        if-exceeding {
            bandwidth-limit 1m;
            burst-size-limit 63k;
        }
    }
}
family inet {
    prefix-action per-source-policer {
        policer 1Mbps-policer;
        subnet-prefix-length 24;
        source-prefix-length 32;
    }
}
filter limit-all-hosts {
    term one {
        from {
            source-address {
                10.10.10/24;
                10.11/16;
            }
        }
    }
}
}

```

```

    }
  }
  then prefix-action per-source-policer;
}
}
}

```

Examples: Classifying Traffic

Classify expedited forwarding traffic:

```

[edit]
firewall {
  policer ef-policer {
    if-exceeding {
      bandwidth-limit 300k;
      burst-size-limit 50k;
    }
    then {
      discard;
    }
  }
  term ef-multifield {
    then {
      loss-priority low;
      forwarding-class expedited-forwarding;
      policer ef-policer;
    }
  }
}

```

Classify assured forwarding traffic:

```

firewall {
  policer af-policer {
    if-exceeding {
      bandwidth-limit 300k;
      burst-size-limit 500k;
    }
    then {
      loss-priority high;
    }
  }
  term af-multifield {
    then {
      loss-priority low;
      forwarding-class assured-forwarding;
      policer af-policer;
    }
  }
}

```

Configuring an Interface Set

In addition to including policers in firewall filters, you can configure an interface set that is not part of a firewall filter configuration. An interface set groups a number of interfaces into one interface set name.

To configure an interface set, include the `interface-set` statement at the `[edit firewall]` hierarchy level:

```
[edit firewall]
interface-set interface-set-name {
  [ interface-names ];
}
```

You must specify more than one interface name to configure an interface set. This interface set can be used for firewall filter matching.

Applying an Interface Policer


In addition to including policers in firewall filters, you can apply an interface policer that is not part of a firewall filter configuration. An interface policer can be applied to each family on an interface.

To apply an interface policer, include the `policer` statement at the `[edit interfaces interface-name unit logical-unit-number family family-name]` hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family family-name]
policer {
  input policer-name;
  output policer-name;
}
```

You must first configure the policer at the `[edit firewall]` hierarchy level before you can apply it to an interface. Both input and output policers are allowed, and can be used in conjunction with existing firewall filters. Input interface policers are evaluated before any input firewall filters. Likewise, output interface policers are evaluated after any output firewall filters (see Figure 12 on page 245).

Figure 12: Incoming and Outgoing Interface Policers



To display a policer on a particular interface, issue the `show interfaces policers` command at the command-line interface (CLI).

Example: Applying an Interface Policer

Apply a policer on circuit cross-connect (CCC) interfaces:

```
[edit interfaces]
so-0/0/0 {
  encapsulation ppp-ccc;
  unit 0 {
    family ccc {
      policer {
        input dragnet;
      }
    }
  }
}
```

Configuring an Aggregate Policer

You can configure a single aggregated policer to limit traffic on the same interface without the use of multiple instances of the same policer. Instead of policing each address family individually on an interface, you can aggregate policing with one policer. This single aggregated policer is also known as the logical-interface policer.

To configure a logical-interface policer, include the **logical-interface-policer** statement at the **[edit firewall policer *policer-name*]** hierarchy level:

```
logical-interface-policer;
```

You can configure rate limiting on the logical-interface policer. For information on configuring rate limiting, see “Configuring Rate Limiting” on page 235. You can configure a policer action for the logical-interface policer. For information on configuring policy actions, see “Configuring a Policer Action” on page 236.

After configuring the aggregated logical-interface policer, you can apply the policer to an interface. To apply an aggregated logical-interface policer, include the **policer *policer-name*** option at the **[edit interfaces *interface-name* unit 0 family *family-name*]** hierarchy level:

```
policer policer-name;
```

For more information about applying policers, see the *JUNOS Class of Service Configuration Guide*.

Example: Configuring an Aggregate Policer

Configure an aggregate policer to rate-limit:

```
[edit firewall policer new-police1]
if-exceeding {
  bandwidth-limit 100m;
  burst-size-limit 500k;
}
logical-interface-policer;
```

```

then {
    discard;
}

```

Apply the aggregate policer to rate-limit IPv4 and IPv6 traffic on interface **fe-0/1/1**:

```

[edit interfaces fe-0/1/1 unit 0 family inet]
policer new-police1;
[edit interfaces fe-0/1/1 unit 0 family inet6]
policer new-police1;

```

Configuring a Bandwidth Policer

The JUNOS software supports policers that rate-limit based on a percentage of physical port speed on an interface.

A bandwidth policer provides similar rate limiting at the logical interface level. For a bandwidth policer, the rate-limiting policer is based on a percentage of the configured logical-interface bandwidth, defined as the shaping rate on that logical interface configured with class-of-service statements.

You can configure a policer to limit the bandwidth and apply that policer on multiple logical interfaces.

To configure a bandwidth policer, include the **logical-bandwidth-policer** statement at the [edit firewall policer *policer-name*] hierarchy level:

```

logical-bandwidth-policer;

```

You can configure rate limiting on the logical-interface policer. For information on configuring rate limiting, see “Configuring Rate Limiting” on page 235. You can configure a policer action for the logical-interface policer. For information on configuring policy actions, see “Configuring a Policer Action” on page 236.

After configuring the bandwidth policer, you can apply the policer to an interface. To apply a bandwidth policer on a logical interface, include the **policer *policer-name*** statement at the [edit interfaces *interface-name* unit 0 family *family-name*] hierarchy level:

```

policer policer-name;

```

For more information about applying policers, see the *JUNOS Class of Service Configuration Guide*.

Example: Configuring a Bandwidth Policer

Configure a bandwidth policer to rate-limit at the logical interface level:

```

[edit firewall policer new-police1]
if-exceeding {
    bandwidth-percent 10;
    burst-size-limit 125k;
}

```

```
logical-bandwidth-policer;
then {
    discard;
}
```

Apply the bandwidth policer to rate-limit IPv4 and IPv6 traffic on interface **fe-0/1/1**:

```
[edit interfaces fe-0/1/1 unit 0 family inet]
policer new-police1;
[edit interfaces fe-0/1/1 unit 0 family inet6]
policer new-police1;
```

Configuring a Load-Balance Group

In addition to including policers in firewall filters, you can configure a load-balance group that is not part of a firewall filter configuration. A load-balance group contains interfaces that all use the same next-hop group characteristic to load-balance the traffic.

To configure a load-balance group, include the **load-balance-group** statement at the **[edit firewall]** hierarchy level:

```
[edit firewall]
load-balance-group group-name {
    next-hop-group [ group-names ];
}
```

Next-hop groups allow you to include multiple interfaces used to forward duplicate packets used in port mirroring. For more information about next-hop groups, see “Configuring a Next-Hop Group” on page 291.

Configuring Tricolor Marking

For T-series and M320 platforms with Enhanced II Flexible PIC Concentrators (FPCs), you can configure single-rate or two-rate tricolor marking (TCM).

TCM extends the functionality of class-of-service (CoS) traffic policing by providing three levels of drop priority instead of two. This allows you to provision more enhanced service-level agreements (SLAs) across the Differentiated Services (DiffServ) domain by defining tricolor marking policers, and three levels of packet loss priority (PLP) for classifiers, rewrite rules, random early detection (RED) drop profiles, and firewall filters.

The color of a packet, as used or set by a tricolor marking policer, corresponds to the packet’s drop precedence (loss priority or PLP). Packets with high PLP are marked red, packets with medium PLP are marked yellow, and packets with low PLP are marked green.

The following sections describe tricolor marking policers:

- Configuring a Tricolor Marking Policer on page 249
- Configuring an Interface Policar Using Tricolor Marking Policing on page 251

Configuring a Tricolor Marking Policer

A tricolor marking policer polices traffic on the basis of metering, including the committed information rate (CIR), the peak information rate (PIR), and their associated burst sizes.

To configure a tricolor marking policer, include the `three-color-policer` statement at the `[edit firewall]` hierarchy level:

```
[edit firewall]
three-color-policer (Configuring) name {
  single-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    peak-information-rate bps;
    peak-burst-size bytes;
  }
}
```

When you configure this type of policer, you can set up to three loss priorities: low, medium-high, and high.



NOTE: To configure a policer that marks packets so that they have medium-low loss priority, you must configure a policer at the `[edit firewall policer policer-name]` hierarchy level.

For example:

```
[edit firewall]
policer 4PLP {
  if-exceeding {
    bandwidth-limit 40k;
    burst-size-limit 4k;
  }
  then loss-priority medium-low;
}
```

Apply this policer at one or both of the following hierarchy levels:

- `[edit firewall family family filter filter-name term rule-name then policer policer-name]`
- `[edit interfaces interface-name unit logical-unit-number family family filter]`

Specify the `single-rate` statement to configure marking based on CIR. If a packet exceeds the CIR in a single-rate policer, it is evaluated by the CBS. Specify the

committed-burst-size option value to configure the maximum number of bytes allowed for incoming packets to burst above the CIR, but still be marked green. Specify the **excess-burst-size** option value to configure the maximum number of bytes allowed for incoming packets to burst above the CIR, but be marked red.

Specify the **two-rate** statement to configure marking based on CIR and PIR. If a packet exceeds the CIR in a two-rate policer, it is evaluated by the PIR. Specify the **committed-information-rate** option value to configure the guaranteed bandwidth under normal line conditions, and the rate up to which packets are marked green. Specify the **committed-burst-size** option value to configure the maximum number of bytes allowed for incoming packets to burst above the CIR, but still be marked green.

Specify the **peak-information-rate** option value to configure the maximum achievable rate. Packets that exceed the CIR, but are below the PIR, are marked yellow. Packets that exceed the PIR are marked red. Specify the **peak-burst-size** option value to configure the maximum number of bytes allowed for incoming packets to burst above the PIR, but still be marked yellow.

For both the **single-rate** statement and the **two-rate** statement, specify the **color-aware** option value to configure metering by preclassification. Metering can increase a PLP, but cannot decrease it. Specify the **color-blind** option value to ignore any preclassification.

For more information about tricolor marking, see the *JUNOS Class of Service Configuration Guide*.

Example: Configuring a Tricolor Marking Policer

Configure a tricolor policer:

```
[edit firewall]
three-color-policer trtcm1 {
  two-rate {
    color-blind;
    committed-information-rate 1048576;
    committed-burst-size 65536;
    peak-information-rate 10485760;
    peak-burst-size 131072;
  }
}
```

Apply the tricolor policer to a firewall filter.

```
[edit firewall]
filter fil {
  term default {
    then {
      three-color-policer {
        two-rate trtcm1;
      }
    }
  }
}
```


Configuring an Interface Policer Using Tricolor Marking Policing

You can configure a policer to limit traffic on an interface in the ingress or egress direction. Instead of policing each address family individually on an interface, you can aggregate policing with one policer. This single aggregated policer is known as the logical-interface policer. You can configure tricolor marking policing to limit the bandwidth through a logical interface.

To configure a policer on a logical interface using tricolor marking policing, include the `action` statement and the `logical-interface-policer` statement at the `[edit firewall three-color-policer name]` hierarchy level:

```
[edit firewall]
three-color-policer (Configuring) policer-name {
  action {
    loss-priority high then discard;
  }
  logical-interface-policer;
  single-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    peak-information-rate bps;
    peak-burst-size bytes;
  }
}
```

For detailed information about bandwidth policers on a logical interface, see “Configuring an Aggregate Policer” on page 246.

You can configure separate policing on the ingress and egress direction on the logical interface.

Example: Rate-Limiting Bandwidth Using Tricolor Marking Policing

Configure tricolor marking policing on a logical interface to rate-limit the bandwidth on the logical interface.

```
[edit firewall]
three-color-policer trtcm-1 {
  action {
    loss-priority high then discard;
  }
  logical-interface-policer;
  two-rate {
    color-blind;
    committed-information-rate 1500000;
    committed-burst-size 150k;
  }
}
```

```

        peak-information-rate 3m;
        peak-burst-size 300k;
    }
}

```

Examples: Configuring Policing

The following example shows a complete filter configuration containing a policer. It limits all FTP traffic from a given source to certain rate limits. Traffic exceeding the limits is discarded, and the remaining traffic is accepted and counted.

```

[edit]
firewall {
  policer policer-1 {
    if-exceeding {
      bandwidth-limit 400k;
      burst-size-limit 100k;
    }
    then {
      discard;
    }
  }
  term tcp-ftp {
    from {
      source-address 10.2.3/24;
      protocol tcp;
      destination-port ftp;
    }
    then {
      policer policer-1;
      accept;
      count count-ftp;
    }
  }
}

```

The following example shows a complete filter configuration containing two policers, and includes the **next term** action. Policer **policer-1** limits all traffic from a given source to certain rate limits, then sets the forwarding class. Policer **policer-2** limits all traffic to a second set of rate limits. Traffic exceeding the limits is discarded; the remaining traffic is accepted.

```

[edit]
firewall {
  policer policer-1 {
    if-exceeding {
      bandwidth-limit 10m;
      burst-size-limit 100k;
    }
    then {
      forwarding-class 0;
    }
  }
  policer policer-2 {
    if-exceeding {

```

```

        bandwidth-limit 100m;
        burst-size-limit 100k;
    }
    then {
        discard;
    }
}
filter f {
    term term-1 {
        then {
            policer policer-1;
            next term;
        }
    }
    term term-2 {
        then {
            policer policer-2;
            accept;
        }
    }
}
}

```

The following example limits all FTP traffic from a given source to certain rate limits, but defines the policer outside the filter, thereby creating a template that can be referenced by more than one filter or more than one term within a filter. Traffic exceeding the limits is discarded, and the remaining traffic is accepted and counted.

```

[edit]
firewall {
    policer policer-1 {
        if-exceeding {
            bandwidth-limit 400k;
            burst-size-limit 100k;
        }
        then {
            discard;
        }
    }
    filter limit-ftp {
        term tcp-ftp {
            from {
                source-address 10.2.3/24;
                protocol tcp;
                destination-port ftp;
            }
            then {
                policer policer-1;
                accept;
                count count-ftp;
            }
        }
    }
}
}

```

The following example shows a filter intended to thwart denial-of-service (DoS) SYN attacks:

```
[edit]
firewall {
  policer syn-recvd {
    if-exceeding {
      bandwidth-limit 40k;
      burst-size-limit 15000;
    }
    then discard;
  }
  term allow-syn {
    from {
      source-address {
        192.168.12.50/32; # trusted addresses
      }
    }
    then {
      log;
      accept;
    }
  }
  term limit-syn {
    from {
      protocol tcp;
      tcp-initial;
    }
    then {
      count limit-syn;
      policer syn-recvd;
      accept;
    }
  }
  term default {
    then accept;
  }
}
[edit] # apply filter to lo0 to control traffic to the Routing Engine
interfaces {
  lo0 {
    unit 0 {
      family inet {
        filter {
          input syn-attack;
        }
      }
      address 172.16.4.53/32;
    }
  }
}
```

The following example uses one filter to do the following:

- Stop all User Datagram Protocol (UDP) and Internet Control Message Protocol (ICMP) traffic destined to these addresses (in **term a**).

- Send ICMP through the policer (in **term b**).
- Accept ICMP traffic within contract and all other traffic (in **term c**).



NOTE: It is important to keep the terms in order; once a packet has a match within the firewall filter, it is not examined in subsequent terms. For example, if you configured the filter to send ICMP traffic through the policer before discarding ICMP and UDP traffic to the addresses (in **term a**), you would not get the desired result.

```
[edit firewall]
policer policer-1 {
  if-exceeding {
    bandwidth-limit 200k;
    burst-size-limit 3k;
  }
  then {
    loss-priority 1;
    forwarding-class 1;
  }
}
term a {
  from {
    destination-address {
      10.126.50.2/23;
      10.130.12.1/23;
      10.82.16.0/24 except;
      10.82.0.3/18;
    }
    protocol [icmp udp];
  }
  then {
    count packets-dropped;
    discard;
  }
}
term b {
  from {
    protocol icmp;
  }
  then policer policer-1;
}
term c {
  then accept;
}
```


Chapter 12

Summary of Firewall Filter and Policer Configuration Statements

The following descriptions explain each of the firewall filter and policer configuration statements. The statements are organized alphabetically.

accounting-profile

Syntax	accounting-profile <i>name</i> ;
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable collection of accounting data for the specified filter.
Options	<i>name</i> —Name assigned to the accounting profile.
Usage Guidelines	See “Configuring a Firewall Filter Accounting Profile” on page 221.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

action

Syntax	action { loss-priority high then discard; }
Hierarchy Level	[edit firewall three-color-policer <i>name</i>], [edit logical-systems <i>logical-system-name</i> firewall three-color-policer <i>name</i>]
Release Information	Statement introduced in JUNOS Release 8.2. Logical systems support introduced in JUNOS Release 9.3.
Description	Discard traffic on a logical interface using tricolor marking policing.
Options	The statements are explained separately in this chapter.
Usage Guidelines	See “Configuring an Interface Policer Using Tricolor Marking Policing” on page 251.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

family

Syntax

```
family family-name {
    filter filter-name {
        accounting-profile name;
        interface-specific;
    }
    prefix-action name {
        count;
        destination-prefix-length prefix-length;
        policer policer-name;
        source-prefix-length prefix-length;
        subnet-prefix-length prefix-length;
    }
    simple-filter filter-name {
        term term-name {
            from {
                match-conditions;
            }
            then {
                action;
                action-modifiers;
            }
        }
    }
}
```

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

Release Information Statement introduced before JUNOS Release 7.4.
Logical systems support introduced in JUNOS Release 9.3.
simple-filter statement introduced in JUNOS Release 7.6.
any family type introduced in JUNOS Release 8.0.
(MX-series routers only) **bridge** family type introduced in JUNOS Release 8.4.

Description Configure a firewall filter for IP version 4 (IPv4) or IP version 6 (IPv6) traffic. On the MX-series only, configure a firewall filter for Layer 2 traffic in a bridging environment.

Options *family-name*—Version or type of addressing protocol:

- **any**—Protocol-independent match conditions.
- **bridge**—(MX-series routers only) Layer 2 packets that are part of bridging domain.
- **ccc**—Layer 2 switching cross-connects.
- **inet**—IPv4 addressing protocol.
- **inet6**—IPv6 addressing protocol.
- **mpls**—Multiprotocol Label Switching (MPLS).
- **vpls**—Virtual private LAN service.

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Family Address Type” on page 166.

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

filter

Syntax

```
filter filter-name {
    accounting-profile name;
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            match-conditions;
        }
        then {
            action;
            action-modifiers;
        }
    }
}
```

Hierarchy Level [edit firewall family *family-name*],
[edit logical-systems *logical-system-name* firewall family *family-name*]

Release Information Statement introduced before JUNOS Release 7.4.
Logical systems support introduced in JUNOS Release 9.3.

Description Configure firewall filters.

Options *filter-name*—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (" ").

The remaining statements are explained separately.

Usage Guidelines See “Configuring Firewall Filters” on page 165.

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

filter-specific

Syntax	filter-specific;
Hierarchy Level	[edit firewall policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4. Logical systems support introduced in JUNOS Release 9.3.
Description	Configure a policer to act as a filter-specific policer. If this statement is not specified, then the policer defaults to a term-specific policer.
Usage Guidelines	See “Configuring Filter-Specific Policers” on page 239.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

firewall

Syntax	firewall { ... }
Hierarchy Level	[edit], [edit logical-systems <i>logical-system-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4. Logical systems support introduced in JUNOS Release 9.3.
Description	Configure firewall filters. The statements are explained separately.
Usage Guidelines	See “Firewall Filter Configuration” on page 163.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.

if-exceeding

Syntax	if-exceeding { bandwidth-limit <i>bps</i> ; bandwidth-percent <i>number</i> ; burst-size-limit <i>bytes</i> ; }
Hierarchy Level	[edit firewall policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4. Logical systems support introduced in JUNOS Release 9.3.
Description	Configure policer rate limits.
Options	<p>bandwidth-limit <i>bps</i>—Traffic rate, in bits per second (bps). There is no minimum value, but any value below 61,040 bps results in an effective rate of 30,520 bps. Range: 32,000 through 32,000,000,000 bps Default: None</p> <p>bandwidth-percent <i>number</i>—Port speed, in decimal percentage number. Range: 1 through 100 Default: None</p> <p>burst-size-limit <i>bytes</i>—Maximum burst size. The minimum recommended value is the maximum transmission unit (MTU) of the IP packets being policed. The maximum burst size limit varies depending on the platform used. Range: 1500 through 100,000,000 bytes Default: None</p>
Usage Guidelines	See “Configuring Rate Limiting” on page 235.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.

interface-set

Syntax	<code>interface-set <i>interface-set-name</i> { [<i>interface-names</i>]; }</code>
Hierarchy Level	[edit firewall], [edit logical-systems <i>logical-system-name</i> firewall]
Release Information	Statement introduced before JUNOS Release 7.4. Logical systems support introduced in JUNOS Release 9.3.
Description	Configure an interface set.
Options	<i>interface-names</i> —Names of interfaces to include in the interface set. You must specify more than one name.
Usage Guidelines	See “Configuring an Interface Set” on page 245.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.

interface-specific

Syntax	<code>interface-specific;</code>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i>], [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> filter <i>filter-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4. Logical systems support introduced in JUNOS Release 9.3.
Description	Configure interface-specific names for firewall counters.
Usage Guidelines	See “Configuring Interface-Specific Counters” on page 206.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

load-balance-group

Syntax	<code>load-balance-group <i>group-name</i> { next-hop-group [<i>group-names</i>]; }</code>
Hierarchy Level	[edit firewall]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure a load-balance group.
Options	<i>group-name</i> —Name of load-balance group. <i>group-names</i> —Name of next-hop groups to include in the load-balance group set.
Usage Guidelines	See “Configuring a Load-Balance Group” on page 248.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.

logical-bandwidth-policer

Syntax	<code>logical-bandwidth-policer;</code>
Hierarchy Level	[edit firewall policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i>]
Release Information	Statement introduced in JUNOS Release 8.2. Logical systems support introduced in JUNOS Release 9.3.
Description	Extend the policer rate limits to logical interfaces. The policer rate limit is based on the shaping rate defined on the logical interface.
Usage Guidelines	See “Configuring a Bandwidth Policer” on page 247.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

logical-interface-policer

Syntax	logical-interface-policer;
Hierarchy Level	[edit firewall policer <i>policer-name</i>], [edit firewall three-color-policer <i>name</i>], [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall three-color-policer <i>name</i>]
Release Information	Statement introduced before JUNOS Release 7.4. Support at the [edit firewall three-color-policer <i>name</i>] hierarchy level introduced in JUNOS Release 8.2. Logical systems support introduced in JUNOS Release 9.3.
Description	Configure an aggregate policer.
Usage Guidelines	See “Configuring an Aggregate Policer” on page 246 and “Configuring Tricolor Marking” on page 248.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

policer

Syntax	<pre> policer <i>policer-name</i> { filter-specific; if-exceeding { bandwidth-limit <i>bps</i>; bandwidth-percent <i>number</i>; burst-size-limit <i>bytes</i>; } logical-interface-policer; then { <i>policer-action</i>; } } </pre>
Hierarchy Level	[edit firewall], [edit logical-systems <i>logical-system-name</i> firewall]
Release Information	Statement introduced before JUNOS Release 7.4. Logical systems support introduced in JUNOS Release 9.3.
Description	Configure policer rate limits and actions. When included at the [edit firewall] hierarchy level, it creates a template, and you do not have to configure a policer individually for every firewall filter or interface. To activate a policer, you must include the policer action modifier in the then statement in a firewall filter term or on an interface.
Options	<p><i>policer-action</i>—One or more actions to take:</p> <ul style="list-style-type: none"> ■ discard—Discard traffic that exceeds the rate limits. ■ forwarding-class <i>class-name</i>—Specify the particular forwarding class. ■ loss-priority—Set the packet loss priority (PLP) to low, medium-low, medium-high, or high. <p><i>policer-name</i>—Name that identifies the policer. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p>then—Actions to take on matching packets.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Policers” on page 235.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.

prefix-action

Syntax	<pre> prefix-action <i>name</i> { count; destination-prefix-length <i>prefix-length</i>; policer <i>policer-name</i>; source-prefix-length <i>prefix-length</i>; subnet-prefix-length <i>prefix-length</i>; }</pre>
Hierarchy Level	[edit firewall family inet], [edit logical-systems <i>logical-system-name</i> firewall family inet]
Release Information	Statement introduced before JUNOS Release 7.4. Logical systems support introduced in JUNOS Release 9.3.
Description	Configure prefix-specific action.
Options	<p>count—Enable counter.</p> <p>destination-prefix-length <i>prefix-length</i>—Destination prefix length. Range: 0 through 32</p> <p>policer <i>policer-name</i>—Policer name.</p> <p>source-prefix-length <i>prefix-length</i>—Source prefix length. Range: 0 through 32</p> <p>subnet-prefix-length <i>prefix-length</i>—Subnet prefix length. Range: 0 through 32</p>
Usage Guidelines	See “Configuring Prefix-Specific Actions” on page 239.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.

service-filter

Syntax	<pre> service-filter <i>filter-name</i> { term <i>term-name</i> { from { <i>match-conditions</i>; } then { <i>action</i>; <i>action-modifiers</i>; } } } </pre>
Hierarchy Level	[edit firewall family inet], [edit logical-systems <i>logical-system-name</i> firewall family inet]
Release Information	Statement introduced before JUNOS Release 7.4. Logical systems support introduced in JUNOS Release 9.3.
Description	Configure service filters.
Options	<p><i>filter-name</i>—Name that identifies the service filter. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Service Filters” on page 203.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration

simple-filter

Syntax	<pre> simple-filter <i>filter-name</i> { term <i>term-name</i> { from { <i>match-conditions</i>; } then { <i>action</i>; <i>action-modifiers</i>; } } } </pre>
Hierarchy Level	[edit firewall family inet], [edit logical-systems <i>logical-system-name</i> firewall family inet]
Release Information	Statement introduced in JUNOS Release 7.6. Logical systems support added in JUNOS Release 9.3.
Description	Configure simple filters.
Options	<p><i>filter-name</i>—Name that identifies the simple filter. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Simple Filters” on page 203.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration

term

Syntax

```
term term-name {
    filter filter-name;
    from {
        match-conditions;
    }
    then {
        actions;
        action-modifiers;
    }
}
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name],
[edit firewall family family-name service-filter filter-name],
[edit firewall family family-name simple-filter filter-name],
[edit logical-systems logical-system-name firewall family family-name filter filter-name],
[edit logical-systems logical-system-name firewall family family-name service-filter
filter-name],
[edit logical-systems logical-system-name firewall family family-name simple-filter
filter-name]
```

Release Information

Statement introduced before JUNOS Release 7.4.
 filter option introduced in JUNOS Release 7.6.
 Logical systems support introduced in JUNOS Release 9.3.

Description

Define a firewall filter term.

Options

actions—(Optional) An action to take if conditions match. If you do not specify an action, the packets that match the conditions in the **from** statement are accepted. The actions are described in Table 28 on page 169.

action-modifiers—(Optional) One or more actions to perform on a packet. The action modifiers are described in Table 28 on page 169.

filter-name—(Optional) A filter within a filter. This term references another filter.

from—(Optional) Match packet fields to values. If not included, all packets are considered to match and the actions and action modifiers in the **then** statement are taken.

match-conditions—One or more conditions to use to make a match. The conditions are described in Table 29 on page 176, Table 32 on page 186, and Table 33 on page 189.

term-name—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (" ").

then—(Optional) Actions to take on matching packets. If not included and a packet matches all the conditions in the **from** statement, the packet is accepted.

Usage Guidelines See “Configuring Firewall Filters” on page 165.

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

three-color-policer

See the following sections:

- three-color-policer (Applying) on page 272
- three-color-policer (Configuring) on page 273

three-color-policer (Applying)

Syntax three-color-policer {
 (single-rate | two-rate) *policer-name*;
 }

Hierarchy Level [edit firewall family *family-name* filter *filter-name* term *term-name* then],
 [edit logical-systems *logical-system-name* firewall family *family-name* filter *filter-name*
 term *term-name* then]

Release Information Statement introduced in JUNOS Release 7.4.
 single-rate statement added in JUNOS Release 8.2.
 Logical systems support introduced in JUNOS Release 9.3.

Description For T-series and M320 platforms with Enhanced II Flexible PIC Concentrators (FPCs)
 and the T640 platform with Enhanced Scaling FPC4, apply a tricolor marking policer.

Options single-rate—Named tricolor policer is a single-rate policer.

two-rate—Named tricolor policer is a two-rate policer.

policer-name—Name of a tricolor policer.

Usage Guidelines See “Configuring a Filter Action Statement” on page 168.

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

three-color-policer (Configuring)

Syntax	<pre> three-color-policer <i>policer-name</i> { action { loss-priority high then discard; } logical-interface-policer; single-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; excess-burst-size <i>bytes</i>; } two-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; peak-information-rate <i>bps</i>; peak-burst-size <i>bytes</i>; } } </pre>
Hierarchy Level	[edit firewall], [edit logical-systems <i>logical-system-name</i> firewall]
Release Information	Statement introduced in JUNOS Release 7.4. action statement introduced in JUNOS Release 8.2. Logical systems support introduced in JUNOS Release 9.3.
Description	Configure a tricolor marking policer.
Options	<p>color-aware—Metering varies by preclassification. Metering can increase a packet's assigned PLP, but cannot decrease it.</p> <p>color-blind—Packet preclassification is ignored. All packets are evaluated by the CBS. If a packet exceeds the CBS, it is evaluated by the EBS.</p> <p>committed-burst-size <i>bytes</i>—Maximum bytes allowed for incoming packets to be marked green. Range: 1500 through 100,000,000,000 bytes</p> <p>committed-information-rate <i>bps</i>—Guaranteed bandwidth under normal line conditions, and the average rate up to which packets are marked green. Range: 32,000 through 40,000,000,000 bps</p> <p>excess-burst-size <i>bytes</i>—Maximum bytes allowed for incoming packets. Packets that exceed the EBS are marked red. Range: 1500 through 100,000,000,000 bytes</p> <p>peak-burst-size <i>bytes</i>—Maximum bytes allowed for incoming packets to burst above the PIR, but still be marked yellow. Range: 1500 through 100,000,000,000 bytes</p>

peak-information-rate *bps*—Maximum achievable rate. Packets that exceed the CIR but are below the PIR are marked yellow. Packets that exceed the PIR are marked red.

Range: 32,000 through 40,000,000,000 bps

single-rate—Marking is based on the CIR, CBS, and the EBS.

two-rate—Marking is based on the CIR and the PIR.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Tricolor Marking” on page 248.

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

virtual-channel

Syntax `virtual-channel virtual-channel-name;`

Hierarchy Level [edit firewall family *family-name* filter *filter-name* term *term-name* then],
 [edit logical-systems *logical-system-name* firewall family *family-name* filter *filter-name* term *term-name* then]

Release Information Statement introduced before JUNOS Release 7.4.
 Logical systems support introduced in JUNOS Release 9.3.

Description For J-series Services Routers only, select the traffic to be transmitted by way of a particular virtual channel. *virtual-channel-name* must be one of the names that you define at the [edit class-of-service virtual-channels] hierarchy level.

Options *virtual-channel-name*—Name of the virtual channel.

Usage Guidelines See the *JUNOS Class of Service Configuration Guide*.

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

Part 4

Traffic Sampling and Forwarding

- Traffic Sampling and Forwarding Overview on page 277
- Traffic Sampling and Forwarding Configuration on page 279
- Summary of Traffic Sampling and Forwarding Options Configuration Statements on page 335

Chapter 13

Traffic Sampling and Forwarding Overview

Traffic sampling allows you to sample IP traffic based on particular input interfaces and various fields in the packet header. You can also use traffic sampling to monitor any combination of specific logical interfaces, specific protocols on one or more interfaces, a range of addresses on a logical interface, or individual IP addresses. Information about the sampled packets is saved to files on the router's hard disk.

The forwarding policies allow you to configure the per-flow load balancing, port mirroring, and Domain Name System (DNS) or Trivial File Transfer Protocol (TFTP) forwarding. Beginning with JUNOS Release 9.0, you can configure per-prefix load balancing. This feature enables the router to elect the next hop independent of the route chosen by other routers. The result is a better utilization of available links.

Traffic sampling and forwarding are supported only on routing platforms equipped with an Internet Processor II application-specific integrated circuit (ASIC). To determine whether a routing platform has an Internet Processor II ASIC, use the **show chassis hardware** command.

Traffic sampling is not meant to capture all packets received by a router. We do not recommend excessive sampling (a rate greater than 1/1000 packets), because it can increase the load on your processor. If you need to set a higher sampling rate to diagnose a particular problem or type of traffic received, we recommend that you revert to a lower sampling rate after you discover the problem or troublesome traffic.

Chapter 14

Traffic Sampling and Forwarding Configuration

To configure forwarding options and traffic sampling, include statements at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
accounting group-name {
  output {
    cflowd [ hostnames ] {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      autonomous-system-type (origin | peer);
      port port-number;
      version format;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface [ interface-names ] {
      engine-id number;
      engine-type number;
      source-address address;
    }
  }
}
dhcp-relay {
  traceoptions {
    flag all;
    flag database;
    flag state;
    flag interface;
    flag rtsock;
    flag packet;
    flag packet-option;
    flag io;
    flag ha;
    flag ui;
```

```

flag general;
flag fwd;
flag rpd;
file file-name {
    <files number>;
    <size maximum-file-size>;
    <match regex>;
    <(world-readable | no-world-readable)>;
}
}
overrides {
    always-write-giaddr;
    always-write-option-82;
    layer2-unicast-replies;
    trust-option-82;
    disable-relay;
}
relay-option-60 {
    vendor-option {
        (equals | starts-with) (ascii match-string | hexadecimal match-hex) {
            (relay-server-group server-group-name | local-server-group
            local-server-group-name | drop);
        }
        (default-relay-server-group server-group-name | default-local-server-group
        local-server-group-name | drop);
    }
}
relay-option-82 {
    circuit-id {
        prefix host-name logical-system-name routing-instance-name;
    }
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
active-server-group server-group-name;
group group-name {
    active-server-group server-group-name;
    overrides {
        always-write-giaddr;
        always-write-option-82;
        layer2-unicast-replies;
        trust-option-82;
        disable-relay;
    }
    relay-option-60 {
        vendor-option {
            (equals | starts-with) (ascii match-string | hexadecimal match-hex) {
                (relay-server-group server-group-name | local-server-group
                local-server-group-name | drop);
            }
            (default-relay-server-group server-group-name | default-local-server-group
            local-server-group-name | drop);
        }
    }
}

```

```

    }
    relay-option-82 {
        circuit-id {
            prefix host-name logical-system-name routing-instance-name;
        }
    }
    [interface interface-name;]
}
}
family family-name {
    filter {
        input filter-name;
        output filter-name;
    }
    route-accounting;
}
flood {
    input filter-name;
}
hash-key {
    family inet {
        layer-3;
        layer-4;
    }
    family mpls {
        no-interface-index;
        label-1;
        label-2;
        label-3;
        no-labels;
        no-label-1-exp;
        payload {
            ip {
                layer-3-only;
                port-data {
                    source-msb;
                    source-lsb;
                    destination-msb;
                    destination-lsb;
                }
            }
        }
    }
}
family multiservice {
    destination-mac;
    source-mac;
}
}
helpers {
    bootp {
        client-response-ttl;
        description description-of-service;
        interface interface-group {
            client-response-ttl number;
            description description-of-interface;
            maximum-hop-count number;
        }
    }
}

```

```

        minimum-wait-time seconds;
        no-listen;
        server [ addresses ];
    }
    maximum-hop-count number;
    minimum-wait-time seconds;
    server [ addresses ];
}
domain {
    description description-of-service;
    server < [ routing-instance routing-instance-names ] >;
    interface interface-name {
        description description-of-interface;
        no-listen;
        server < [ routing-instance routing-instance-names ] >;
    }
}
tftp {
    description description-of-service;
    server < [ routing-instance routing-instance-names ] >;
    interface interface-name {
        description description-of-interface;
        no-listen;
        server < [ routing-instance routing-instance-names ] >;
    }
}
traceoptions {
    file {
        files;
        size;
    }
    flag;
    level;
}
}
load-balance {
    indexed-next-hop;
    per-flow {
        hash-seed number;
    }
    per-prefix {
        hash-seed number;
    }
}
}
monitoring group-name {
    family inet {
        output {
            cflowd hostname {
                port port-number;
            }
            export-format cflowd-version-5;
            flow-active-timeout seconds;
            flow-export-destination {
                collector-pic;
            }
            flow-inactive-timeout seconds;

```



```

        interface interface-name {
            engine-id number;
            engine-type number;
            input-interface-index number;
            output-interface-index number;
            source-address address;
        }
    }
}
next-hop-group [ group-names ] {
    interface interface-name {
        next-hop [ addresses ];
    }
}
packet-capture {
    disable;
    file {
        filename file-name;
        files number;
        size number;
        (world-readable | no-world-readable);
    }
    maximum-capture-size number;
}
port-mirroring {
    family (inet | inet6) {
        input {
            rate number;
            run-length number;
        }
        output {
            interface interface-name {
                next-hop address;
            }
            no-filter-check;
        }
    }
}
traceoptions {
    file filename {
        files number;
        size bytes;
        (world-readable | no-world-readable);
    }
}
sampling {
    disable;
    input {
        family (inet | mpls) {
            max-packets-per-second number;
            rate number;
            run-length number;
        }
    }
    output {
        cflowd hostname {

```

```

    version9 {
        template template-name;
    }
    aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
            caida-compliant;
        }
        source-prefix;
    }
    autonomous-system-type (origin | peer);
    (local-dump | no-local-dump);
    port port-number;
    source-address address;
    version format;
}
file {
    disable;
    filename filename;
    files number;
    size bytes;
    (stamp | no-stamp);
    (world-readable | no-world-readable);
}
flow-active-timeout seconds;
flow-inactive-timeout seconds;
interface [ interface-names ] {
    engine-id number;
    engine-type number;
    source-address address;
}
}
traceoptions {
    file filename {
        files number;
        size bytes;
        (world-readable | no-world-readable);
    }
}
}

```

This chapter describes the following tasks for configuring traffic sampling and forwarding options:

- Minimum Traffic Sampling or Forwarding Configuration on page 285
- Configuring a Forwarding Table Filter on page 286
- Configuring IPv6 Accounting on page 287
- Configuring Traffic Sampling on page 288
- Configuring Discard Accounting on page 289
- Configuring Flow Monitoring on page 290
- Configuring a Next-Hop Group on page 291

- Configuring Per-Flow Load-Balancing Information on page 292
- Configuring Per-Prefix Load Balancing on page 293
- Configuring Per-Flow Load Balancing Based on Hash Values on page 293
- Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent on page 294
- Configuring DNS and TFTP Packet Forwarding on page 295
- Configuring the Extended DHCP Relay Agent on page 299
- Disabling Traffic Sampling on page 321
- Examples: Configuring Traffic Sampling on page 321
- Configuring Traffic Sampling Output on page 324
- Tracing Traffic Sampling Operations on page 326
- Configuring Flow Aggregation (cflowd) on page 326
- Configuring Active Flow Monitoring Using Version 9 on page 329
- Configuring Port Mirroring on page 330
- Configuring Packet Capture on page 332

Minimum Traffic Sampling or Forwarding Configuration

To configure traffic sampling, you must perform at least the following tasks:

1. Create a firewall filter to apply to the logical interfaces being sampled by including the `filter` statement at the `[edit firewall family family-name]` hierarchy level. In the filter then statement, you must specify the action modifier `sample` and the action `accept`.

```
[edit firewall family]
filter filter-name {
  term term-name {
    then {
      sample;
      accept;
    }
  }
}
```

2. Apply the filter to the interfaces on which you want to sample traffic:

```
[edit interfaces]
interface-name {
  unit logical-unit-number {
    family family-name {
      filter {
        input filter-name;
      }
      address address {
        destination destination-address;
      }
    }
  }
}
```

```
    }
  }
```

3. Enable **sampling** and specify a nonzero sampling rate:

```
[edit forwarding-options]
sampling {
  input {
    family inet {
      rate number;
    }
  }
}
```

Configuring a Forwarding Table Filter

A forwarding table filter allows you to filter data packets based on their components and perform an action on packets that match the filter. You can filter on the ingress or egress packets of a forwarding table.

To apply a forwarding table filter at the ingress of a forwarding table, include the **filter input** statement at the `[edit forwarding-options family family-name]` hierarchy level:

```
[edit forwarding-options family family-name]
filter {
  input filter-name;
}
```

On the MX-series router only, to apply a forwarding table filter for a virtual switch, include the **filter input** statement at the `[edit routing-instances routing-instance-name bridge-domains bridge-domain-name forwarding-options]` hierarchy level:

```
[edit routing-instances routing-instance-name {
  bridge-domains bridge-domain-name {
    forwarding-options {
      filter {
        input filter-name;
      }
    }
  }
}
```

For more information about how to configure a virtual switch, see the *JUNOS MX-series Layer 2 Configuration Guide*.

You can filter based upon destination-class information by configuring a firewall filter on the egress of the forwarding table. By applying firewall filters to packets that have been forwarded by a routing table, you can match based on certain parameters that are decided by the route lookup. For example, routes can be classified into specific destination and source classes. Firewall filters used for policing and mirroring are able to match based upon these classes.

To apply a firewall filter at the egress of a forwarding table, include the **filter output** statement at the `[edit forwarding-options family family-name]` hierarchy level:

```
[edit forwarding-options family family-name]
filter {
  output filter-name;
}
```



NOTE: The egress forwarding table filter will be applied on the ingress of the flexible PIC concentrator (FPC). If different packets to the same destination arrive on different FPCs, they may encounter different policers.



NOTE: You cannot configure both an egress forwarding table filter and the `interface-group` statement at the `[edit interfaces family inet filter]` hierarchy level. The egress forwarding table filter is applied to transit packets only.



NOTE: The egress forwarding table filter is not supported for the J-series Services Routers.

To apply a forwarding table filter to a flood table, include the `flood input` statement at the `[edit forwarding-options family family-name]` hierarchy level:

```
[edit forwarding-options family family-name]
flood {
  input filter-name;
}
```



NOTE: The flood statement is valid for the `vpls` address family only.

Configuring IPv6 Accounting

You can configure the routing platform to track IPv6 specific packets and bytes passing through the router.

To enable IPv6 accounting, include the `route-accounting` statement at the `[edit forwarding-options family inet6]` hierarchy level:

```
[edit forwarding-options family inet6]
route-accounting;
```

By default, IPv6 accounting is disabled. If IPv6 accounting is enabled, it is disabled after a reboot of the routing platform. To view IPv6 statistics, issue the `show interface statistics operational` mode command. For more information, see the *JUNOS Interfaces Command Reference*.

Configuring Traffic Sampling

On routing platforms containing a Monitoring Services Physical Interface Card (PIC) or an Adaptive Services PIC, you can configure traffic sampling for traffic passing through the routing platform.

To configure traffic sampling on a logical interface, include the `input` statement at the `[edit forwarding-options sampling]` hierarchy level:

```
[edit forwarding-options sampling]
input {
  family (inet | mpls) {
    max-packets-per-second number;
    rate number;
    run-length number;
  }
}
```

Beginning with JUNOS Release 8.3, you can also configure traffic sampling of MPLS traffic.

Specify the threshold traffic value by using the `max-packets-per-second` statement. The value is the maximum number of packets to be sampled, beyond which the sampling mechanism begins dropping packets. The range is 0 through 65,535. A value of 0 instructs the Packet Forwarding Engine not to sample any packets. The default value is 1000.



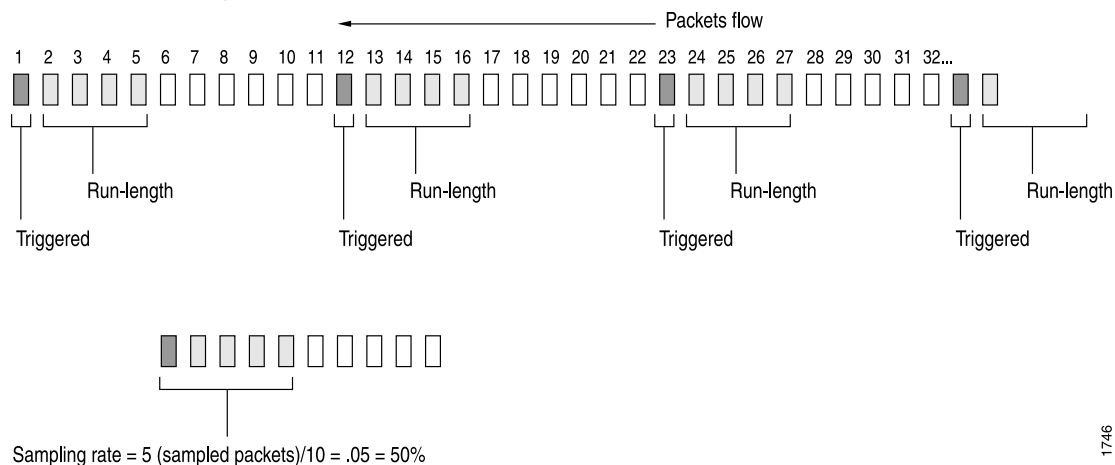
NOTE: This statement is not valid for port mirroring.

Specify the sampling rate by setting the values for `rate` and `run-length` (see Figure 13 on page 288).

Figure 13: Configure Sampling Rate

Rate and Run-length

Case #1 Rate = 10, run-length = 4



The **rate** statement specifies the ratio of packets to be sampled. For example, if you configure a rate of 10, x number of packets out of every 10 is sampled, where $x = \text{run-length} + 1$. By default, the rate is 0, which means that no traffic is sampled.

The **run-length** statement specifies the number of matching packets to sample following the initial one-packet trigger event. Configuring a run length greater than 0 allows you to sample packets following those already being sampled.

If you do not include the **input** statement, sampling is disabled.

To collect the sampled packets in a file, include the **file** statement at the [edit forwarding-options sampling output] hierarchy level. For more information about the output file formats, see “Configuring Traffic Sampling Output” on page 324.

You can also send the sampled packets to a specified host using the cflowd version 5 and 8 formats or the version 9 format as defined in RFC 3954. For more information, see “Configuring Flow Aggregation (cflowd)” on page 326 and “Configuring Active Flow Monitoring Using Version 9” on page 329.

The JUNOS software does not sample packets originating from the router. If you configure a sampling filter and apply it to the output side of an interface, then only the transit packets going through that interface are sampled. Packets that are sent from the Routing Engine to the Packet Forwarding Engine are not sampled.

When you apply a firewall filter to a loopback interface, the filter might block responses from the Monitoring Services PIC. To allow responses from the Monitoring Services PIC to pass through for sampling purposes, configure a term in the firewall filter to include the Monitoring Services PIC’s IP address. For more detailed information about configuring firewall filters, see “Firewall Filter Configuration” on page 163.

Configuring Discard Accounting

On routing platforms containing a Monitoring Services PIC or an Adaptive Services PIC, you can configure accounting for traffic passing through the routing platform.

To configure discard accounting, include the **accounting** statement at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
accounting group-name {
  output {
    cflowd [ hostnames ] {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
    }
    autonomous-system-type (origin | peer);
  }
}
```

```

        port port-number;
        version format;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface [ interface-names ] {
        engine-id number;
        engine-type number;
        source-address address;
    }
}

```

To configure an accounting group, include the **accounting** statement and specify a *group-name*. To configure the output flow aggregation, include the **cflowd** statement. For more information about flow aggregation, see “Configuring Flow Aggregation (cflowd)” on page 326. To configure the interval before exporting an active flow, include the **flow-active-timeout** statement. The default value for **flow-active-timeout** is 1800 seconds. To configure the interval before a flow is considered inactive, include the **flow-inactive-timeout** statement. The default value for **flow-inactive-timeout** is 60 seconds. To configure the interface that sends out monitored information, include the **interface** statement. Discard accounting is supported for the Monitoring Services PIC only.

When you apply a firewall filter to a loopback interface, the filter might block responses from the Monitoring Services PIC. To allow responses from the Monitoring Services PIC to pass through for accounting purposes, configure a term in the firewall filter to include the Monitoring Services PIC IP address. For more detailed information about configuring firewall filters, see “Firewall Filter Configuration” on page 163.

You can use discard accounting for passive and active flow monitoring. For more detailed information about configuring passive and active flow monitoring, see the *JUNOS Feature Guide* and the *JUNOS Class of Service Configuration Guide*.

Configuring Flow Monitoring

On routing platforms containing the Monitoring Services PIC or the Monitoring Services II PIC, you can configure flow monitoring for traffic passing through the routing platform. This type of monitoring method is passive monitoring.

To configure flow monitoring, include the **monitoring** statement at the [edit forwarding-options hierarchy level:

```

[edit forwarding-options]
monitoring group-name {
    family inet {
        output {
            cflowd hostname {
                port port-number;
            }
        }
        export-format cflowd-version-5;
        flow-active-timeout seconds;
        flow-export-destination {
            (cflowd-collector | collector-pic);
        }
    }
}

```



```

    }
    flow-inactive-timeout seconds;
    interface interface-name {
        engine-id number;
        engine-type number;
        input-interface-index number;
        output-interface-index number;
        source-address address;
    }
}
}
}

```

To configure a passive monitoring group, include the **monitoring** statement and specify a group name. To configure monitoring on a specified address family, include the **family** statement and specify an address family. To specify an interface to monitor incoming traffic, include the **input** statement. To configure the monitoring information that is sent out, include the **output** statement. To configure the output flow aggregation, include the **cflowd** statement. For more information about flow aggregation, see “Configuring Flow Aggregation (cflowd)” on page 326. To specify the format of the monitoring information sent out, include the **export-format** statement and specify a version number. To configure the interval before exporting an active flow, include the **flow-active-timeout** statement. The default value for **flow-active-timeout** is 1800 seconds. To enable flow collection, include the **flow-export-destination** statement. To configure the interval before a flow is considered inactive, include the **flow-inactive-timeout** statement. The default value for **flow-inactive-timeout** is 60 seconds. To configure the interface that sends out the monitored information, include the **interface** statement. Flow monitoring is supported for Monitoring Services PIC interfaces only.

When you apply a firewall filter to a loopback interface, the filter might block responses from the Monitoring Services PIC. To allow responses from the Monitoring Services PIC to pass through for monitoring purposes, configure a term in the firewall filter to include the Monitoring Services PIC’s IP address. For more detailed information about configuring firewall filters, see “Firewall Filter Configuration” on page 163.

For more detailed information about configuring passive and active flow monitoring, see the *JUNOS Feature Guide* and the *JUNOS Class of Service Configuration Guide*.

Configuring a Next-Hop Group

Next-hop groups allow you to include multiple interfaces used to forward duplicate packets used in port mirroring.

To configure a next-hop group, include the **next-hop-group** statement at the [edit forwarding-options] hierarchy level:

```

[edit forwarding-options]
next-hop-group [ group-names ] {
    interface interface-name {
        next-hop [ addresses ];
    }
}

```

```
}
}
```

You can specify one or more group names. To configure the interface that sends out sampled information, include the **interface** statement and specify an interface. To specify a next-hop address to send sampled information, include the **next-hop** statement and specify an IP address.

Next-hop groups have the following restrictions:

- Next-hop groups are supported for M-series routers only.
- Next-hop groups support up to 16 next-hop addresses.
- You can configure up to 30 next-hop groups.
- Each next-hop group must have at least two next-hop addresses.



NOTE: When routes are exported, RIPv2 supports third-party next hops specified in policies, such as Virtual Router Redundancy Protocol (VRRP) groups.

Next-hop groups can be used for port mirroring. For more information about configuring port mirroring, see “Configuring Port Mirroring” on page 330 and the *JUNOS Feature Guide*.

Configuring Per-Flow Load-Balancing Information

By default, when there are multiple equal-cost paths to the same destination, the JUNOS software chooses one of the next-hop addresses at random.

On routing platforms with the Internet Processor II application-specific integrated circuit (ASIC), you have three additional options:

- You can specify what information the routing platform uses for per-flow load balancing based on port data (instead of on source and destination IP addresses only). For aggregated Ethernet and aggregated SONET/SDH interfaces, you can load-balance based on the Multiprotocol Label Switching (MPLS) label information. For more information, see “Configuring the Load-Balance Per-Packet Action” on page 136.
- You can also configure per-prefix load balancing, which allows you to configure a hash value that enables the router to elect a next hop independently of the route chosen by other routers. For more information, see “Configuring Per-Prefix Load Balancing” on page 293.
- You can also configure per-flow load balancing, which allows you to configure a unique, load balance hash value for each PFE slot. For more information, see “Configuring Per-Flow Load Balancing Based on Hash Values” on page 293.

Configuring Per-Prefix Load Balancing

By default, the JUNOS software uses a hashing method based only on the destination address to elect a forwarding next hop when multiple equal-cost paths are available. As a result, when multiple routers share the same set of forwarding next hops for a given destination, they can elect the same forwarding next hop.

Beginning with JUNOS Release 9.0, you can enable router-specific load balancing by including a per-prefix hash value. However, this method applies only to indirect next hops. In other words, when we have a route with a protocol next hop that is not directly connected, it can be resolved over a set of equal-cost forwarding next hops. Only in this case, we use the hashing algorithm to elect a forwarding next hop. An example of this is routes learned from an IBGP neighbor. The protocol next hop for those routes might not be directly reachable and would be resolved through some IGP or static routes. The result could be a set of equal-cost forwarding next hops to reach that protocol next hop. Per-prefix load balancing thus leads to better utilization of the available links.

To configure per-prefix load balancing, include the **load-balance** statement at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
load-balance {
  indexed-next-hop;
  per-prefix {
    hash-seed number;
  }
}
```

To enable per-prefix load balancing, you must include the **hash-seed *number*** statement. The range that you can configure is 0 through 65,535. 0 is the default value; if no hash seed is configured, the elected forwarding next hop would be the same as in previous releases.

To generate a permuted index of next-hop entries for unicast and aggregate next hops, include the **indexed-next-hop** statement at the [edit forwarding-options load-balance] hierarchy level:

```
indexed-next-hop;
```

Configuring Per-Flow Load Balancing Based on Hash Values

By default, the JUNOS software uses a hashing method based only on the destination address to elect a forwarding next hop when multiple equal-cost paths are available. All PFE slots are assigned the same hash value by default.

Beginning with JUNOS Release 9.3, you can enable router-specific load balancing by configuring a unique, load balance hash value for each PFE slot.

To configure per-prefix load balancing, include the **load-balance** statement at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
```

```

load-balance {
  indexed-next-hop;
  per-flow {
    hash-seed number;
  }
}

```

To enable per-flow load balancing, you must include the `hash-seed number` statement. The range that you can configure is 0 through 65,535. 0 is the default value; if no hash seed is configured, the elected forwarding next hop would be the same as in previous releases.

Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent

You can configure the router or an interface to act as a Dynamic Host Configuration Protocol (DHCP) or Bootstrap Protocol (BOOTP) relay agent. This means that a locally attached host can issue a DHCP or BOOTP request as a broadcast message. If the router or an interface sees this broadcast message, it relays the message to a specified DHCP or BOOTP server.

You should configure the router or an interface to be a DHCP/BOOTP relay agent if you have locally attached hosts and a distant DHCP or BOOTP server.

To configure the router to act as a DHCP/BOOTP relay agent, include the `bootp` statement at the `[edit forwarding-options helpers]` hierarchy level:

```

[edit forwarding-options helpers]
bootp {
  client-response-ttl number;
  description description-of-service;
  interface interface-group {
    client-response-ttl number;
    description description-of-interface;
    maximum-hop-count number;
    minimum-wait-time seconds;
    no-listen;
    server [ ];
  }
  maximum-hop-count number;
  minimum-wait-time seconds;
  server [ ];
}

```

To set the description of the BOOTP service, DHCP service, or interface, include the `description` statement.

To set a logical interface or a group of logical interfaces with a specific DHCP-relay or BOOTP configuration, include the `interface` statement.

To set the routing instance of the server to forward, include the `routing-instance` statement. You can include as many routing instances as necessary in the same statement.

To stop packets from being forwarded on a logical interface, a group of logical interfaces, or the router, include the `no-listen` statement.

To set the maximum allowed number in the hops field of the BOOTP header, include the `maximum-hop-count` statement. Headers that have a larger number in the hops field are not forwarded. If you omit the `maximum-hop-count` statement, the default value is 4 hops.

To set the minimum allowed number of seconds in the `secs` field of the BOOTP header, include the `minimum-wait-time` statement. Headers that have a smaller number in the `secs` field are not forwarded. The default value for the minimum wait time is zero (0).

To set the IP address or addresses that specify the DHCP or BOOTP server for the router or interface, include the `server` statement. You can include as many addresses as necessary in the same statement.

To perform an IP time-to-live (TTL) decrement before accepting any packets for processing when DHCP response packets are sent to a DHCP client, include the `client-response-ttl` statement.

You can also configure an individual logical interface to be a DHCP/BOOTP relay if you have locally attached hosts and a remote DHCP or BOOTP server connected to one of the router's interfaces. For more information, see the *JUNOS System Basics Configuration Guide*.

Configuring DNS and TFTP Packet Forwarding

You can configure the router to support Domain Name System (DNS) and Trivial File Transfer Protocol (TFTP) packet forwarding for IPv4 traffic, which allows clients to send DNS or TFTP requests to the router. The responding DNS or TFTP server recognizes the client address and sends a response directly to that address. By default, the router ignores DNS and TFTP request packets.

To enable DNS or TFTP packet forwarding, include the `helpers` statement at the `[edit forwarding-options]` hierarchy level:

```
[edit forwarding-options]
helpers {
  domain {
    description description-of-service;
    server < [ routing-instance routing-instance-names ] >;
    interface interface-name {
      description description-of-interface;
      no-listen;
      server < [ routing-instance routing-instance-names ] >;
    }
  }
  tftp {
    description description-of-service;
    server < [ routing-instance routing-instance-names ] >;
    interface interface-name {
      description description-of-interface;
    }
  }
}
```

```

        no-listen;
        server < [ routing-instance routing-instance-names ] >;
    }
}

```

To set domain packet forwarding, include the **domain** statement.

To set the description of the DNS or TFTP service, include the **description** statement.

To set TFTP packet forwarding, include the **tftp** statement.

To set a DNS or TFTP server (with an IPv4 address), include the **server** statement. Use one address for either a global configuration or for each interface.

To set the routing instance of the server to forward, include the **routing-instance** statement. You can include as many routing instances as necessary in the same statement.

To disable recognition of DNS or TFTP requests on one or more interfaces, include the **no-listen** statement. If you do not specify at least one interface with this statement, the forwarding service is global to all interfaces on the routing platform.

The following sections discuss the following topics:

- Tracing BOOTP, DNS, and TFTP Forwarding Operations on page 296
- Example: Configuring DNS Packet Forwarding on page 299

Tracing BOOTP, DNS, and TFTP Forwarding Operations

BOOTP, DNS, and TFTP forwarding tracing operations track all BOOTP, DNS, and TFTP operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

By default, nothing is traced. If you include the **traceoptions** statement at the [edit forwarding-options helpers] hierarchy level, the default tracing behavior is the following:

- Important events are logged in a file called **fud** located in the **/var/log** directory.
- When the file **fud** reaches 128 kilobytes (KB), it is renamed **fud.0**, then **fud.1**, and so on, until there are 3 trace files. Then the oldest trace file (**fud.2**) is overwritten. (For more information about how log files are created, see the *JUNOS System Log Messages Reference*.)
- Log files can be accessed only by the user who configures the tracing operation.

You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements at the [edit forwarding-options helpers] hierarchy level:

```

[edit forwarding-options helpers]
file filename <files number> <size size> <(world-readable | no-world-readable)> <match
  regex>;
flag {

```

```

all;
configuration;
database;
events;
policy;
server;
timer-events;
}

```

These statements are described in the following sections:

- Configuring the Log Filename on page 297
- Configuring the Number and Size of Log Files on page 297
- Configuring Access to the Log File on page 297
- Configuring a Regular Expression for Lines to Be Logged on page 298
- Configuring the Trace Operations on page 298

Configuring the Log Filename

By default, the name of the file that records trace output is `fud`. You can specify a different name by including the `file` statement at the `[edit event-options traceoptions]` hierarchy level:

```

[edit forwarding-options helpers]
file filename;

```

Configuring the Number and Size of Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed `filename.0`, then `filename.1`, and so on, until there are three trace files. Then the oldest trace file (`filename.2`) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the `[edit forwarding-options helpers]` hierarchy level:

```

[edit forwarding-options helpers]
file files number size size;

```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (`filename`) reaches 2 MB, `filename` is renamed `filename.0`, and a new file called `filename` is created. When the new `filename` reaches 2 MB, `filename.0` is renamed `filename.1` and `filename` is renamed `filename.0`. This process repeats until there are 20 trace files. Then the oldest file (`filename.19`) is overwritten by the newest file (`filename.0`).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

Configuring Access to the Log File

By default, log files can be accessed only by the user who configures the tracing operation.

To specify that any user can read all log files, include the `file world-readable` statement at the `[edit forwarding-options helpers]` hierarchy level:

```
[edit forwarding-options helpers]
file world-readable;
```

To explicitly set the default behavior, include the `file no-world-readable` statement at the `[edit forwarding-options helpers]` hierarchy level:

```
[edit forwarding-options helpers]
file no-world-readable;
```

Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including the `match` statement at the `[edit forwarding-options helpers file filename]` hierarchy level and specifying a regular expression (regex) to be matched:

```
[edit forwarding-options helpers]
file filename match regex;
```

Configuring the Trace Operations

By default, only important events are logged. You can configure the trace operations to be logged by including the following statements at the `[edit forwarding-options helpers traceoptions]` hierarchy level:

```
[edit forwarding-options helpers traceoptions]
flag {
  all;
  configuration;
  events;
  database;
  policy;
  server;
  timer-events;
}
```

Table 37 on page 298 describes the meaning of the tracing flags.

Table 37: Tracing Flags

Flag	Description	Default Setting
all	Trace all operations.	Off
configuration	Log reading of configuration at the <code>[edit forwarding-options helpers]</code> hierarchy level.	Off
events	Trace important events.	Off

Table 37: Tracing Flags *(continued)*

Flag	Description	Default Setting
database	Log events involving storage and retrieval in events database.	Off
policy	Log policy processing.	Off
server	Log communication with processes that are generating events.	Off
timer-events	Log internally generated events.	Off

To display the end of the log, issue the `show log (fud | last)` operational mode command:

```
[edit]
user@host# run show log (fud | last)
```

Example: Configuring DNS Packet Forwarding

Enable DNS packet request forwarding to all interfaces on the router except `t1-1/1/2` and `t1-1/1/3`:

```
[edit forwarding-options helpers]
dns {
  server 10.10.10.30;
  interface {
    t1-1/1/2 {
      no-listen;
      server 10.10.10.9;
    }
    t1-1/1/3 {
      no-listen;
      server 10.10.10.4;
    }
  }
}
```

Configuring the Extended DHCP Relay Agent

You can configure extended DHCP relay options on the router and enable the router to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server. You can use DHCP relay in carrier edge applications such as video/IPTV to obtain configuration parameters, including an IP address, for your subscribers.

To configure the extended DHCP relay agent on the router, include the `dhcp-relay` statement:

```
dhcp-relay {
  traceoptions {
```

```

    flag all;
    flag auth;
    flag database;
    flag state;
    flag interface;
    flag rtsock;
    flag packet;
    flag packet-option;
    flag io;
    flag ha;
    flag ui;
    flag general;
    flag fwd;
    flag rpd;
    flag session-db;
    file file-name {
        <files number>;
        <size maximum-file-size>;
        <match regex>;
        <(world-readable | no-world-readable)>;
    }
}
overrides {
    always-write-giaddr;
    always-write-option-82;
    layer2-unicast-replies;
    trust-option-82;
    disable-relay;
}
relay-option-60 {
    vendor-option {
        (equals | starts-with) (ascii match-string | hexadecimal match-hex) {
            (relay-server-group server-group-name | local-server-group
            local-server-group-name | drop);
        }
        (default-relay-server-group server-group-name | default-local-server-group
        local-server-group-name | drop);
    }
}
relay-option-82 {
    circuit-id {
        prefix host-name logical-system-name routing-instance-name;
    }
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
active-server-group server-group-name;
authentication {
    password password-string;
    username-include {
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
    }
}

```

```

        logical-system-name;
        mac-address;
        option-60;
        option-82 [circuit-id] [remote-id];
        routing-instance-name;
        user-prefix user-prefix-string;
    }
}
group group-name {
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            logical-system-name;
            mac-address;
            option-60;
            option-82 [circuit-id] [remote-id];
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    overrides {
        always-write-giaddr;
        always-write-option-82;
        layer2-unicast-replies;
        trust-option-82;
        disable-relay;
    }
    relay-option-60 {
        vendor-option {
            (equals | starts-with) (ascii match-string | hexadecimal match-hex) {
                (relay-server-group server-group-name | local-server-group
                 local-server-group-name | drop);
            }
            (default-relay-server-group server-group-name | default-local-server-group
             local-server-group-name | drop);
        }
    }
    relay-option-82 {
        circuit-id {
            prefix host-name logical-system-name routing-instance-name;
        }
    }
    [interface [upto upto-interface-name] [exclude];
}
}

```

You can include these statements at the following hierarchy levels:

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

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options]
- [edit routing-instances *routing-instance-name* forwarding-options]

For more information about how to use the DHCP relay agent in a video/IPTV application, see the *JUNOS Feature Guide*.



NOTE: The extended DHCP relay agent options configured with the **dhcp-relay** statement are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, you cannot enable both the extended DHCP relay agent and the DHCP/BOOTP relay agent on the router at the same time.

For information about the DHCP/BOOTP relay agent, see “Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent” on page 294.

The following sections describe the operation and configuration of the extended DHCP relay agent and provide simple and complex configuration examples:

- Interaction Between the DHCP Relay Agent, DHCP Client, and DHCP Servers on page 303
- Access and Access-Internal Routes on page 303
- DHCP State Persistence on page 304
- Graceful Routing Engine Switchover on page 304
- Overriding the Default DHCP Relay Configuration on page 304
- Using Option 60 Information to Forward Client Traffic to Specific DHCP Servers on page 306
- Enabling and Disabling Insertion of Option 82 Information on page 308
- Configuring Server Groups on page 311
- Configuring Active Server Groups on page 311
- Grouping Interfaces with Common DHCP Relay Configuration on page 312
- Using External AAA Authentication Services on page 313
- Verifying and Managing DHCP Relay Agent Clients on page 314
- Tracing Extended DHCP Relay Agent Operations on page 314
- Example: Minimum DHCP Relay Agent Configuration on page 317
- Example: DHCP Relay Agent Configuration with Multiple Clients and Servers on page 318
- Example: Using Option 60 Strings to Forward DHCP Client Traffic on page 319
- Example: Using Option 60 Strings to Drop DHCP Client Traffic on page 320

Interaction Between the DHCP Relay Agent, DHCP Client, and DHCP Servers

In a typical carrier edge network configuration, the DHCP client is on the subscriber's computer, and the DHCP relay agent is configured on the router between the DHCP client and one or more DHCP servers. The following steps describe, at a high level, how the DHCP client, DHCP relay agent, and DHCP server interact in a configuration that includes two DHCP servers.

1. The DHCP client sends a discover packet to find a DHCP server in the network from which to obtain configuration parameters for the subscriber, including an IP address.
2. The DHCP relay agent receives the discover packet and forwards copies to each of the two DHCP servers. The DHCP relay agent then creates an entry in its internal client table to keep track of the client's state.
3. In response to receiving the discover packet, each DHCP server sends an offer packet to the client. The DHCP relay agent receives the offer packets and forwards them to the DHCP client.
4. On receipt of the offer packets, the DHCP client selects the DHCP server from which to obtain configuration information. Typically, the client selects the server that offers the longest lease time on the IP address.
5. The DHCP client sends a request packet that specifies the DHCP server from which to obtain configuration information.
6. The DHCP relay agent receives the request packet and forwards copies to each of the two DHCP servers.
7. The DHCP server requested by the client sends an acknowledgement (ACK) packet that contains the client's configuration parameters.
8. The DHCP relay agent receives the ACK packet and forwards it to the client.
9. The DHCP client receives the ACK packet and stores the configuration information.
10. If configured to do so, the DHCP relay agent installs a host route and Address Resolution Protocol (ARP) entry for this client.
11. After establishing the initial lease on the IP address, the DHCP client and the DHCP server use unicast transmission to negotiate lease renewal or release. The DHCP relay agent "snoops" on all of the packets unicast between the client and the server that pass through the router to determine when the lease for this client has expired or been released. This process is referred to as lease shadowing or passive snooping.

Access and Access-Internal Routes

The DHCP application on a video services router uses both access routes and access-internal routes to represent either the end users or the networks behind the attached router. An access route represents a network behind an attached video services router, and is set to a preference of 13. An access-internal route is a /32 route that represents a directly attached end user, and is set to a preference of 12.

To configure import and export of access routes and access-internal routes in a routing policy, include the `access` and `access-internal` keywords as match conditions at the `[edit policy-options policy-statement policy-name term term-name from protocol]` hierarchy level. For information, see “Importing and Exporting Access and Access-Internal Routes in a Routing Policy” on page 57.

To display configuration information for access routes and access-internal routes, use the `show route extensive`, `show route protocol access`, and `show route protocol access-internal` operational commands. For command syntax and examples, see the *JUNOS Routing Protocols and Policies Command Reference*.

DHCP State Persistence

The extended DHCP relay agent maintains the state of active DHCP client leases in persistent storage on the router. It can recover this state if the DHCP relay agent process fails or is manually restarted, or if you manually reboot (gracefully shut down) the router. DHCP state persistence prevents the loss of active DHCP clients in either of these circumstances. If a power failure occurs or if the kernel stops operating on a single Routing Engine, however, the state of active DHCP client leases is lost.

DHCP state persistence is automatically enabled when you configure the extended DHCP relay agent on the router by including the `dhcp-relay` statement.

The DHCP relay agent records in persistent storage only those DHCP clients that are fully bound, which means that they currently have an active lease on an IP address from a DHCP server. DHCP clients in a renewal or rebind state are considered to be fully bound, and their state is also maintained in persistent storage. When a DHCP client lease expires or the client is released, the DHCP relay agent removes the client state from persistent storage.

Graceful Routing Engine Switchover

The extended DHCP relay agent supports graceful Routing Engine switchover on all routing platforms that contain dual Routing Engines. To support graceful Routing Engine switchover, the DHCP relay agent automatically mirrors (replicates) information about the state of bound DHCP clients from the master Routing Engine to the backup Routing Engine.

To enable graceful Routing Engine switchover support for the extended DHCP relay agent, include the `graceful-switchover` statement at the `[edit chassis redundancy]` hierarchy level. You cannot disable graceful Routing Engine switchover support for the extended DHCP relay agent when the router is configured to support graceful Routing Engine switchover.

For more information about using graceful Routing Engine switchover, see the *JUNOS High Availability Configuration Guide*.

Overriding the Default DHCP Relay Configuration

To override the default configuration settings for the extended DHCP relay agent, include the `overrides` statement:

```

overrides {
    always-write-giaddr;
    always-write-option-82;
    layer2-unicast-replies;
    trust-option-82;
    disable-relay;
}

```

To override global DHCP relay agent configuration options, include the **overrides** statement and its subordinate statements at the **[edit forwarding-options dhcp-relay]** hierarchy level. To override DHCP relay agent configuration options for a named group of interfaces, include the statements at the **[edit forwarding-options dhcp-relay group *group-name*]** hierarchy level.

To remove all DHCP relay agent configuration overrides at a particular hierarchy level, include the **overrides** statement with no subordinate statements.

Overwriting giaddr Information

You can configure the DHCP relay agent to change the gateway IP address (giaddr) field in packets that it forwards between a DHCP client and a DHCP server.

To overwrite the giaddr of every DHCP packet with the giaddr of the DHCP relay agent before forwarding the packet to the DHCP server, include the **always-write-giaddr** statement.

Overriding Option 82 Information

You can configure the DHCP relay agent to add or remove the DHCP relay agent information option (option 82) in DHCP packets.

To override the default option 82 information in DHCP packets destined for a DHCP server, include the **always-write-option-82** statement. The use of this statement causes the DHCP relay agent to perform one of the following actions, depending on how it is configured:

- If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server.
- If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.

Using Layer 2 Unicast Transmission for DHCP Packets

To override the setting of the broadcast bit in DHCP request packets and instead use the layer 2 unicast transmission method to send DHCP Offer reply packets and DHCP ACK reply packets from the DHCP server to DHCP clients during the discovery process, include the **layer2-unicast-replies** statement.

Trusting Option 82 Information

By default, the DHCP relay agent treats client packets with a giaddr of 0 (zero) and option 82 information as if they originated at an untrusted source, and drops them without further processing. To override this behavior and instead enable the DHCP relay agent to process DHCP client packets that have a giaddr of 0 (zero) and contain option 82 information, include the `trust-option-82` statement.

Disabling DHCP Relay

To disable DHCP relay on specific interfaces in a group, include the `disable-relay` statement.

Using Option 60 Information to Forward Client Traffic to Specific DHCP Servers

You can configure the extended DHCP relay agent to use the DHCP vendor class identifier option (option 60) in DHCP client packets to forward client traffic to specific DHCP servers. This feature is useful in network environments where DHCP clients access services provided by multiple vendors and DHCP servers. For example, a DHCP client might gain Internet access from a particular DHCP server provided by one vendor, and access IPTV service from a different DHCP server provided by another vendor. The option 60 string enables vendors to include vendor-specific information in DHCP client packets.

To use option 60 vendor-specific information to select a DHCP server to which to forward the client packets, include the `relay-option-60` statement:

```
relay-option-60 {
  vendor-option {
    (equals | starts-with) (ascii match-string | hexadecimal match-hex) {
      (relay-server-group server-group-name |
       local-server-group local-server-group-name |
       drop);
    }
    (default-relay-server-group server-group-name |
     default-local-server-group local-server-group-name |
     drop);
  }
}
```

To configure option 60 support globally, include the `relay-option-60` statement and subordinate statements at the `[edit forwarding-options dhcp-relay]` hierarchy level. To configure option 60 support for a named group of interfaces, include the `relay-option-60` statement and subordinate statements at the `[edit forwarding-options dhcp-relay group group-name]` hierarchy level. You can also configure option 60 support for the extended DHCP relay agent on a per logical system and per routing instance basis.

Using Matching Option 60 Strings to Process DHCP Client Traffic

Configuring option 60 support helps you manage multivendor networks by enabling the extended DHCP relay agent to compare option 60 vendor-specific strings received

in DHCP client packets against a list of ASCII or hexadecimal strings that you configure on the router.

The match criteria you configure for the option 60 string-to-DHCP server mapping can be either of the following:

- To specify an exact, left-to-right match of the configured match string with the option 60 string, use the **vendor-option equals** statement with either the **ascii** statement (to define a nonempty ASCII match string of 1 through 255 alphanumeric characters) or the **hexadecimal** statement (to define a hexadecimal match string of 1 through 255 hexadecimal characters [0 through 9, a through f, A through F]).
- To specify a partial match of the configured match string with the option 60 string, use the **vendor-option starts-with** statement with either the **ascii** statement or the **hexadecimal** statement. In this case, the option 60 string can contain a superset of the configured ASCII or hexadecimal string, provided that the leftmost characters of the option 60 string entirely match the characters in the configured match string. When you use the **starts-with** statement, the longest match rule applies. For example, the extended DHCP relay agent matches the string “test123” before it matches the string “test”.

If the option 60 string received in the DHCP client packet matches the configured ASCII or hexadecimal string, you can define one of the following actions for the associated DHCP client packets:

- To relay client traffic to a group of specific DHCP relay servers that provide the requested client service, use the **relay-server-group** statement.

The DHCP client packet is relayed to all of the servers specified in the **server-group** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level that map to the vendor class identifier information provided in the option 60 string. To configure a named group of DHCP relay servers, which are also referred to as vendor-option servers, include the **server-group** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level, as described in “Configuring Server Groups” on page 311.

- To forward client traffic to a specific extended DHCP local server, use the **local-server-group** statement.

To configure an extended DHCP local server, include the **dhcp-local-server** statement at the **[edit system services]** hierarchy level. For information about configuring and using the extended DHCP local server, see the *JUNOS System Basics Configuration Guide*.

- To drop (discard) the packets, use the **drop** statement.

Specifying that certain DHCP client packets be dropped can be useful when DHCP clients request services that are invalid or no longer supported.

The following additional considerations apply when you configure an ASCII or hexadecimal match string:

- You can configure the same ASCII or hexadecimal match string as both an exact (**equals**) match and as a partial (**starts-with**) match. In that case, the exact string match configured with the **equals** statement takes precedence over the partial string match configured with the **starts-with** statement.
- A server group can contain multiple server addresses and can map to more than one match string.
- You can configure an unlimited number of match strings.
- The use of wildcard attributes in match strings is not supported.

For configuration examples that illustrate how to use matching option 60 strings to forward or drop DHCP client traffic, see “Example: Using Option 60 Strings to Forward DHCP Client Traffic” on page 319 and “Example: Using Option 60 Strings to Drop DHCP Client Traffic” on page 320.

Using Nonmatching Option 60 Strings to Process DHCP Client Traffic

If the option 60 string received in the DHCP client packet does not match the configured ASCII or hexadecimal string, you can define one of the following default actions for the associated DHCP client packets:

- To relay client traffic to a default extended DHCP relay server that you specify, use the **default-relay-server-group** statement.
- To forward client traffic to a default extended DHCP local server that you specify, use the **default-local-server-group** statement.
- To drop (discard) the packets, use the **drop** statement.

In rare instances, the extended DHCP relay agent might receive a DHCP client packet with an option 60 string of zero (0) length. In this case, there is nothing in the option 60 string against which to match. As a result, such packets are treated as if they contained nonmatching option 60 strings; that is, they can be relayed to a default DHCP relay server, forwarded to a default DHCP extended local server, or dropped.

For configuration examples that illustrate how to use nonmatching option 60 strings to forward or drop DHCP client traffic, see “Example: Using Option 60 Strings to Forward DHCP Client Traffic” on page 319 and “Example: Using Option 60 Strings to Drop DHCP Client Traffic” on page 320.

Displaying a Count of Discarded DHCP Packets with Option 60 Information

To display the number of discarded DHCP client packets containing option 60 vendor-specific information, use the **show dhcp relay statistics** operational command. For information about using this command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Enabling and Disabling Insertion of Option 82 Information

To enable or disable insertion of the DHCP relay agent information option (option 82) in packets destined for a DHCP server, include the **relay-option-82** statement:

```

relay-option-82 {
  circuit-id {
    prefix host-name logical-system-name routing-instance-name;
  }
}

```

To control insertion of option 82 information globally, include the **relay-option-82** statement and subordinate statements at the **[edit forwarding-options dhcp-relay]** hierarchy level. To control insertion of option 82 information for a named group of interfaces, include the **relay-option-82** statement and subordinate statements at the **[edit forwarding-options dhcp-relay group group-name]** hierarchy level.

To restore the default behavior (option 82 information is not inserted into DHCP packets), include the **relay-option-82** statement with no subordinate statements.

Configuring Agent-Circuit-Id Information

If you use the **relay-option-82** statement to enable insertion of option 82 information in DHCP packets, you must also specify at least the **circuit-id** statement to include the agent-circuit-id suboption (suboption 1) of the DHCP relay agent information option.

If you specify the **circuit-id** statement, the format of the agent-circuit id information for Fast Ethernet (**fe**) or Gigabit Ethernet (**ge**) interfaces is one of the following, depending on your network configuration:

- For Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual local area networks (VLANs) or stacked VLANs (S-VLANs):

(fe | ge)-fpc/pic/port

- For Fast Ethernet or Gigabit Ethernet interfaces that use VLANs:

(fe | ge)-fpc/pic/port:vlan-id

- For Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs:

(fe | ge)-fpc/pic/port:svlan-id-vlan-id

For example, the following is the agent-circuit-id format for a Gigabit Ethernet interface on Flexible PIC Concentrator (FPC) 4, PIC 1, port 1 with S-VLAN ID 122 and VLAN ID 500:

ge-4/1/1:122-500

Configuring an Option 82 Prefix

Optionally, you can also include the **prefix** statement to add a prefix to the base option 82 information in DHCP packets destined for a DHCP server.

```

prefix host-name logical-system-name routing-instance-name;

```

The prefix is separated from the option 82 agent-circuit-id information by a colon (:), and can include any combination of the **host-name**, **logical-system-name**, and

routing-instance-name options. The DHCP relay agent obtains the values for the **host-name**, **logical-system-name**, and **routing-instance-name** as follows:

- If you include the **host-name** option, the DHCP relay agent uses the hostname of the router configured with the **host-name** statement at the **[edit system]** hierarchy level.
- If you include the **logical-system-name** option, the DHCP relay agent uses the logical system name configured with the **logical-system** statement at the **[edit logical-system]** hierarchy level.
- If you include the **routing-instance-name** option, the DHCP relay agent uses the routing instance name configured with the **routing-instance** statement at the **[edit routing-instances]** hierarchy level or at the **[edit logical-system logical-system-name routing-instances]** hierarchy level.

If you include the hostname and either or both the logical system name and the routing instance name in the prefix, the hostname is followed by a forward slash (/). If you include both the logical system name and the routing instance name in the prefix, these values are separated by a semicolon (;).

The following examples show several possible formats for the agent-circuit-id information when you specify the **prefix** statement for Fast Ethernet (**fe**) or Gigabit Ethernet (**ge**) interfaces with S-VLANs.

- If you include only the hostname in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

hostname:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

- If you include only the logical system name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

logical-system-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

- If you include only the routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

- If you include both the hostname and the logical system name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

host-name/logical-system-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

- If you include both the logical system name and the routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

logical-system-name;routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

- If you include the hostname, logical system name, and routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

host-name/logical-system-name;routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

For example, the following is the agent-circuit-id format for a Gigabit Ethernet interface on FPC 4, PIC 1, port 1 with S-VLAN ID 122 and VLAN ID 500. In this example, the prefix consists of the hostname **router1**, the logical system name **xyzcorp**, and the routing instance name **west**.

```
router1/xyzcorp;west:ge-4/1/1:122-500
```

For Fast Ethernet or Gigabit Ethernet interfaces that use VLANs but not S-VLANs, only the *vlan-id* value appears in the agent-circuit-id format. For Fast Ethernet or Gigabit Ethernet interfaces that do not use VLANs or S-VLANs, neither the *vlan-id* value nor the *svlan-id* value appears.

Configuring Server Groups

To configure a named group of DHCP server addresses for use by the extended DHCP relay agent on the router, include the **server-group** statement at the [edit forwarding-options dhcp-relay] hierarchy level:

```
[edit forwarding-options dhcp-relay]
server-group {
  server-group-name {
    server-ip-address;
  }
}
```

You must specify the name of the DHCP server group and the IP addresses of one or more DHCP servers that belong to this group. You can configure a maximum of five IP addresses per named server group.

Configuring Active Server Groups

Configuring an active server group enables you to apply a common DHCP relay agent configuration to a named group of DHCP server addresses. To create an active server group, include the **active-server-group** statement:

```
active-server-group server-group-name;
```

To create an active server group as a global DHCP relay agent configuration option, include the **active-server-group** statement at the [edit forwarding-options dhcp-relay] hierarchy level. To have the group apply only to a named group of interfaces, include the **active-server-group** statement at the [edit forwarding-options dhcp-relay dhcp-relay group group-name] hierarchy level.

Including the **active-server-group** statement at the [edit forwarding-options dhcp-relay dhcp-relay group group-name] hierarchy level (as a group-specific option) overrides the effect of including the **active-server-group** statement at the [edit forwarding-options dhcp-relay] hierarchy level as a global option.

Grouping Interfaces with Common DHCP Relay Configuration

Configuring a DHCP relay group enables you to apply a common DHCP relay agent configuration to a named group of interfaces. To configure a group, include the `group` statement at the `[edit forwarding-options dhcp-relay]` hierarchy level:

```
[edit forwarding-options dhcp-relay]
group group-name {
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      logical-system-name;
      mac-address;
      option-60;
      option-82 [circuit-id] [remote-id];
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  overrides {
    always-write-giaddr;
    always-write-option-82;
    layer2-unicast-replies;
    trust-option-82;
    disable-relay;
  }
  relay-option-60 {
    vendor-option {
      (equals | starts-with) (ascii match-string | hexadecimal match-hex) {
        (relay-server-group server-group-name | local-server-group
          local-server-group-name | drop);
      }
      (default-relay-server-group server-group-name | default-local-server-group
        local-server-group-name | drop);
    }
  }
  relay-option-82 {
    circuit-id {
      prefix host-name logical-system-name routing-instance-name;
    }
  }
  interface [upto upto-interface-name] [exclude];
}
```

Configuring Group-Specific DHCP Relay Options

You can include the following statements at both the `[edit forwarding-options dhcp-relay group group-name]` hierarchy level to set group-specific DHCP relay agent configuration options, and at the `[edit forwarding-options dhcp-relay]` hierarchy level to set global DHCP relay agent configuration options:

- **active-server-group**—Configure an active server group to apply a common DHCP relay agent configuration to a named group of DHCP server addresses. For information, see “Configuring Active Server Groups” on page 311.
- **authentication**—Configure the parameters the router sends to the external AAA server.
- **overrides**—Override the default configuration settings for the extended DHCP relay agent. For information, see “Overriding the Default DHCP Relay Configuration” on page 304.
- **relay-option-60**—Use the DHCP vendor class identifier option (option 60) in DHCP client packets to select a DHCP server to which to forward packets. For more information, see “Using Option 60 Information to Forward Client Traffic to Specific DHCP Servers” on page 306.
- **relay-option-82**—Enable or disable the insertion of option 82 information in packets destined for a DHCP server. For information, see “Enabling and Disabling Insertion of Option 82 Information” on page 308.

The statements configured at the [edit forwarding-options dhcp-relay group *group-name*] hierarchy level apply only to the named group of interfaces, and override any global DHCP relay agent settings configured with the same statements at the [edit dhcp-relay dhcp-relay] hierarchy level.

Enabling the DHCP Relay Agent on Specified Interfaces

To specify the names of one or more interfaces within a specified group on which the DHCP relay agent is enabled, include the **interface** statement at the [edit forwarding-options dhcp-relay group *group-name*] hierarchy level:

```
[edit forwarding-options dhcp-relay group group-name]  
interface interface-name;
```

You can include multiple **interface** *interface-name* statements to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group.

Using External AAA Authentication Services

Both the extended DHCP local server and the extended DHCP relay agent support the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients. When the extended DHCP local server or relay agent receives a discover PDU from a client, the extended DHCP application contacts the AAA server to authenticate the DHCP client. The extended DHCP application can obtain client addresses and DHCP configuration options from the external AAA authentication server.

The external authentication feature also supports AAA directed logout. If the external AAA service supports a user log out directive, the extended DHCP relay agent honors the logout and views it as if it was requested by a CLI management command. All of the client state information and allocated resources are deleted at logout. The extended DHCP relay agent supports directed logout using the list of configured

authentication servers you specify with the **authentication** statement at the [edit access profile *profile-name*] hierarchy.

You can configure either global authentication support or group-specific support.

You must configure the **username-include** statement to enable the use of authentication. The **password** statement is not required and does not cause DHCP to use authentication if the **username-include** statement is not included.

See “Using External AAA Authentication Services” in the *JUNOS System Basics Configuration Guide* for details about configuring external AAA authentication support for the DHCP relay agent.

Verifying and Managing DHCP Relay Agent Clients

To display address bindings or statistics for extended DHCP relay agent clients on the router, use the following operational commands:

- **show dhcp relay binding**
- **show dhcp relay statistics**

To clear address bindings or statistics for DHCP clients, use the following operational commands:

- **clear dhcp relay binding**
- **clear dhcp relay statistics**

For information about using these operational commands, see the *JUNOS Routing Protocols and Policies Command Reference*.

Tracing Extended DHCP Relay Agent Operations

The extended DHCP tracing operations track all extended DHCP relay agent operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

By default, nothing is traced. If you include the **traceoptions** statement at the [edit forwarding-options dhcp-relay] hierarchy level, the default tracing behavior is the following:

- Important events are logged in a file called **jdhcpd** located in the **/var/log** directory. You cannot change the directory (**/var/log**) in which trace files are located.
- When the file **jdhcpd** reaches 128 kilobytes (KB), it is renamed **jdhcpd.0**, then **jdhcpd.1**, and so on, until there are three trace files. Then the oldest trace file (**jdhcpd.2**) is overwritten. (For more information about how log files are created, see the *JUNOS System Log Messages Reference*.)
- Log files can be accessed only by the user who configures the tracing operation.

To trace the extended DHCP relay agent operations, include the **traceoptions** statement at the [edit forwarding-options dhcp-relay] hierarchy level:


```
[edit forwarding-options dhcp-relay]
traceoptions {
  flag all;
  flag auth;
  flag database;
  flag state;
  flag interface;
  flag rtsock;
  flag packet;
  flag packet-option;
  flag io;
  flag ha;
  flag ui;
  flag general;
  flag fwd;
  flag rpd;
  flag session-db;
  file file-name {
    <files number>;
    <size maximum-file-size>;
    <match regex>;
    <(world-readable | no-world-readable)>;
  }
}
```

The extended DHCP relay agent traceoptions statements are described in the following sections:

- Configuring the Extended DHCP Relay Agent Log Filename on page 315
- Configuring the Number and Size of Extended DHCP Relay Agent Log Files on page 315
- Configuring Access to the Extended DHCP Relay Agent Log File on page 316
- Configuring a Regular Expression for Extended DHCP Relay Agent Lines to Be Logged on page 316
- Configuring the Extended DHCP Relay Agent Tracing Flags on page 316

Configuring the Extended DHCP Relay Agent Log Filename

By default, the name of the file that records trace output is `jdhcpd`. You can specify a different name by including the `file` statement at the `[edit forwarding-options dhcp-relay traceoptions]` hierarchy level:

```
file filename;
```

Configuring the Number and Size of Extended DHCP Relay Agent Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed `filename.0`, then `filename.1`, and so on, until there are three trace files. Then the oldest trace file (`filename.2`) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the [edit forwarding-options dhcp-relay traceoptions] hierarchy level:

```
file file-name files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (*filename*) reaches 2 MB, *filename* is renamed *filename.0*, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename.0* is renamed *filename.1* and *filename* is renamed *filename.0*. This process repeats until there are 20 trace files. Then the oldest file (*filename.19*) is overwritten by the newest file (*filename.0*).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

Configuring Access to the Extended DHCP Relay Agent Log File

By default, log files can be accessed only by the user who configures the tracing operation.

To specify that any user can read all log files, include the `file world-readable` statement at the [edit forwarding-options dhcp-relay traceoptions] hierarchy level:

```
file file-name world-readable;
```

To explicitly set the default behavior, include the `file no-world-readable` statement at the [edit forwarding-options dhcp-relay traceoptions] hierarchy level:

```
file file-name no-world-readable;
```

Configuring a Regular Expression for Extended DHCP Relay Agent Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including the `file match` statement at the [edit forwarding-options dhcp-relay traceoptions] hierarchy level and specifying a regular expression (*regex*) to be matched:

```
file file-name match regex;
```

Configuring the Extended DHCP Relay Agent Tracing Flags

By default, only important events are logged. You can configure the trace operations to be logged by including the `flag` statement at the [edit forwarding-options dhcp-relay traceoptions] hierarchy level:

```
flag flag;
```

where *flag* is one of the following extended DHCP relay agent tracing flags:

- **all**—Trace all events
- **auth**—Trace authentication events
- **database**—Trace database events
- **fwd**—Trace firewall process events
- **general**—Trace miscellaneous events
- **ha**—Trace high availability-related events
- **interface**—Trace interface operations
- **io**—Trace I/O operations
- **packet**—Trace packet decoding operations
- **packet-option**—Trace DHCP option decoding operations
- **rpd**—Trace routing protocol process events
- **rtsock**—Trace routing socket operations
- **session-db**—Trace session database events
- **state**—Trace changes in state
- **ui**—Trace user interface operations

To display the end of the log, issue the **show log (jdhcpd | last)** operational mode command:

```
[edit]
user@host# run show log (jdhcpd | last)
```

Example: Minimum DHCP Relay Agent Configuration

The following example shows the minimum configuration you need to use the extended DHCP relay agent on the router:

```
[edit forwarding-options]
dhcp-relay {
  server-group {
    test 10.0.2.1;
  }
  active-server-group test;
  group all {
    interface fe-0/0/2.0;
  }
}
```

This example creates a server group and an active server group named **test** with IP address 10.0.2.1. The DHCP relay agent configuration is applied to a group named **all**. Within this group, the DHCP relay agent is enabled on interface fe-0/0/2.0.

Example: DHCP Relay Agent Configuration with Multiple Clients and Servers

The following example shows a more complex extended DHCP relay agent configuration for a network that includes multiple DHCP clients and DHCP servers. A more detailed explanation follows the example.

```
[edit forwarding-options]
dhcp-relay {
  server-group {
    sp-1 {
      10.0.2.1;
      10.0.2.2;
    }
    sp-2 {
      10.33.2.1;
      10.33.2.2;
      10.33.2.3;
    }
  }
  active-server-group sp-1;
  overrides layer2-unicast-replies;
  group clients_a {
    relay-option-82 circuit-id;
    interface fe-1/0/1.1;
    interface fe-1/0/1.2;
    interface fe-1/0/1.3;
  }
  group clients_b {
    relay-option-82 {
      circuit-id {
        prefix routing-instance-name;
      }
    }
    interface fe-1/0/1.4;
    interface fe-1/0/1.5;
    interface fe-1/0/1.6;
  }
  group eth_dslam_relay {
    active-server-group sp-2;
    overrides {
      trust-option-82;
      layer2-unicast-replies;
    }
    interface fe-1/0/1.7;
    interface fe-1/0/1.8;
    interface fe-1/0/1.9;
  }
}
```

This example creates two server-groups: **sp-1**, which includes DHCP server addresses 10.0.2.1 and 10.0.2.2, and **sp-2**, which includes DHCP server addresses 10.33.2.1, 10.33.2.2, and 10.33.2.3. The active server group to which the DHCP relay agent configuration applies is **sp-1**. A global override is set that causes the DHCP relay agent to use layer 2 unicast transmission to send DHCP reply packets from the DHCP server to DHCP clients during the discovery process.

The example also creates three groups of subscribers and their associated Fast Ethernet interfaces: `clients_a`, `clients_b`, and `eth_dslam_relay`. These groups are configured to meet different needs, as follows:

- The `clients_a` and `clients_b` groups consist of basic subscribers. The service provider for these groups inserts option 82 information in the DHCP packets that are destined for the DHCP server.
- The subscribers in `eth_dslam_relay` are connected to an Ethernet digital subscriber line access multiplexer (DSLAM) that functions as a layer 2 DHCP relay agent. The active server group for `eth_dslam_relay` is `sp-2`. Overrides are set for the `eth_dslam_relay` group that enable the DHCP relay agent to trust option 82 information and to use layer 2 unicast transmission to send DHCP reply packets to DHCP clients during discovery.

Example: Using Option 60 Strings to Forward DHCP Client Traffic

The following extended DHCP relay agent configuration shows how to use the option 60 vendor-specific information in DHCP client packets to forward client traffic to specific DHCP servers. A more detailed explanation follows the example.

```
[edit forwarding-options]
dhcp-relay {
  server-group {
    sp-1 {
      10.0.2.1;
    }
    sp-2 {
      10.33.2.1;
    }
    sp-3 {
      10.22.2.1;
    }
    sp-4 {
      10.10.2.1;
    }
  }
  active-server-group sp-1;
  relay-option-60 {
    vendor-option {
      equals {
        ascii motorola {
          relay-server-group sp-2;
        }
      }
      starts-with {
        hexadecimal ff {
          relay-server-group sp-3;
        }
      }
      default-relay-server-group sp-4;
    }
  }
  group all {
    interface fe-0/0/2.0;
```

```
    }
}
```

This example defines the following actions for DHCP client packets containing option 60 information:

- All packets that contain an exact match with the ASCII string “motorola” are relayed to server group **sp-2**.
- All packets that start with the hexadecimal string “ff” are relayed to server group **sp-3**.
- All packets that do not either exactly match the ASCII string “motorola” or start with the hexadecimal string “ff” are relayed to the default relay server group, **sp-4**.

DHCP client packets that do not contain option 60 information are relayed to the currently configured active server group, **sp-1**.

Server groups **sp-1**, **sp-2**, **sp-3**, and **sp-4** in this example are configured with the `server-group` statement at the `[edit forwarding-options dhcp-relay]` hierarchy level.

Example: Using Option 60 Strings to Drop DHCP Client Traffic

The following extended DHCP relay agent configuration shows how to use the option 60 vendor-specific information in DHCP client packets to drop client traffic. Specifying that certain DHCP client packets be dropped can be useful when DHCP clients request services that are invalid or no longer supported.

```
[edit forwarding-options]
dhcp-relay {
  server-group {
    sp-1 {
      10.0.2.1;
    }
  }
  active-server-group sp-1;
  relay-option-60 {
    vendor-option {
      drop;
    }
  }
  group all {
    interface fe-0/0/2.0;
  }
}
```

In this example, all DHCP client packets containing option 60 information are discarded (dropped), and all packets that do not contain option 60 information are relayed to the currently configured active server group, **sp-1**.

Disabling Traffic Sampling

To explicitly disable traffic sampling on the router, include the `disable` statement at the `[edit forwarding-options sampling]` hierarchy level:

```
[edit forwarding-options sampling]
disable;
```

Examples: Configuring Traffic Sampling

The following sections provide examples of configuring traffic sampling:

- Sampling a Single SONET/SDH Interface on page 321
- Sampling All Traffic from a Single IP Address on page 322
- Sampling All FTP Traffic on page 323

Sampling a Single SONET/SDH Interface

The following configuration gathers statistical sampling information from a small percentage of all traffic on a single SONET/SDH interface and collects it in a file named `sonet-samples.txt`.

Create the filter:

```
[edit firewall family inet]
filter {
  sample-sonet {
    then {
      sample;
      accept;
    }
  }
}
```

Apply the filter to the SONET/SDH interface:

```
[edit interfaces]
so-0/0/1 {
  unit 0 {
    family inet {
      filter {
        input sample-sonet;
      }
      address 10.127.68.254/32 {
        destination 10.127.74.7;
      }
    }
  }
}
```

Finally, configure traffic sampling:

```
[edit forwarding-options]
```

```
sampling {
  input {
    family inet {
      rate 100;
      run-length 2;
    }
  }
  output {
    file {
      filename sonet-samples-txt;
      files 40;
      size 5m;
    }
  }
}
```

Sampling All Traffic from a Single IP Address

The following configuration gathers statistical information about every packet entering the router on a specific Gigabit Ethernet port originating from a single source IP address of 10.45.92.31, and collects it in a file named `samples-10-45-92-31.txt`.

Create the filter:

```
[edit firewall family inet]
filter one-ip {
  term get-ip {
    from {
      source-address 10.45.92.31;
    }
    then {
      sample;
      accept;
    }
  }
}
```

Apply the filter to the Gigabit Ethernet interface:

```
[edit interfaces]
ge-4/1/1 {
  unit 0 {
    family inet {
      filter {
        input one-ip;
      }
      address 10.45.92.254;
    }
  }
}
```

Finally, gather statistics on all the candidate samples; in this case, gather all statistics:

```
[edit forwarding-options]
sampling {
```



```

input {
  family inet {
    rate 1;
  }
}
output {
  file {
    filename samples-215-45-92-31.txt;
    files 100;
    size 100k;
  }
}
}

```

Sampling All FTP Traffic

The following configuration gathers statistical information about a moderate percentage of packets using FTP in the output path of a specific T3 interface, and collects the information in a file named `t3-ftp-traffic.txt`.

Create a filter:

```

[edit firewall family inet]
filter ftp-stats {
  term ftp-usage {
    from {
      destination-port [ftp ftp-data];
    }
    then {
      sample;
      accept;
    }
  }
}
}

```

Apply the filter to the T3 interface:

```

[edit interfaces]
t3-7/0/2 {
  unit 0 {
    family inet {
      filter {
        input ftp-stats;
      }
      address 10.35.78.254/32 {
        destination 10.35.78.4;
      }
    }
  }
}
}

```

Finally, gather statistics on 10 percent of the candidate samples:

```

[edit forwarding-options]
sampling {

```

```

input {
    family inet {
        rate 10;
    }
}
output {
    file {
        filename t3-ftp-traffic.txt;
        files 50;
        size 1m;
    }
}
}

```

Configuring Traffic Sampling Output

You configure traffic sampling results to a file in the `/var/tmp` directory. To collect the sampled packets in a file, include the `file` statement at the `[edit forwarding-options sampling output]` hierarchy level:

```

[edit forwarding-options sampling output]
file {
    disable;
    filename filename;
    files number;
    size bytes;
    (stamp | no-stamp);
    (world-readable | no-world-readable);
}

```

To configure the period of time before an active flow is exported, include the `flow-active-timeout` statement at the `[edit forwarding-options sampling output]` hierarchy level:

```

[edit forwarding-options sampling output]
flow-active-timeout seconds;

```

To configure the period of time before a flow is considered inactive, include the `flow-inactive-timeout` statement at the `[edit forwarding-options sampling output]` hierarchy level:

```

[edit forwarding-options sampling output]
flow-inactive-timeout seconds;

```

To configure the interface, include the `interface` statement at the `[edit forwarding-options sampling output]` hierarchy level:

```

[edit forwarding-options sampling output]
interface [ interface-names ] {
    engine-id number;
    engine-type number;
    source-address address;
}

```

To configure the interval before exporting an active flow, include the **flow-active-timeout** statement. To configure the interval before a flow is considered inactive, include the **flow-inactive-timeout** statement. To configure the interface that sends out monitored information, include the **interface** statement.



NOTE: This feature is not supported with the version 9 template format. You must send traffic flows collected using version 9 to a server. For more information see “Configuring Active Flow Monitoring Using Version 9” on page 329.

Traffic Sampling Output Files

Traffic sampling output is saved to an ASCII text file. The following is an example of the traffic sampling output that is saved to a file in the `/var/tmp` directory. Each line in the output file contains information for one sampled packet. You can optionally display a timestamp for each line.

The column headers are repeated after each group of 1000 packets.

```
# Apr  7 15:48:50
Time                Dest                Src Dest Src Proto TOS Pkt Intf  IP    TCP
                  addr                addr port port          len num frag flags
Apr 7 15:48:54 192.168.9.194 192.168.9.195  0   0   1  0x0  84  8  0x0  0x0
Apr 7 15:48:55 192.168.9.194 192.168.9.195  0   0   1  0x0  84  8  0x0  0x0
Apr 7 15:48:56 192.168.9.194 192.168.9.195  0   0   1  0x0  84  8  0x0  0x0
Apr 7 15:48:57 192.168.9.194 192.168.9.195  0   0   1  0x0  84  8  0x0  0x0
Apr 7 15:48:58 192.168.9.194 192.168.9.195  0   0   1  0x0  84  8  0x0  0x0
```

The output contains the following fields:

- **Time**—Time at which the packet was received (displayed only if you include the **stamp** statement in the configuration)
- **Dest addr**—Destination IP address in the packet
- **Src addr**—Source IP address in the packet
- **Dest port**—Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port for the destination address
- **Src port**—TCP or UDP port for the source address
- **Proto**—Packet’s protocol type
- **TOS**—Contents of the type-of-service (ToS) field in the IP header
- **Pkt len**—Length of the sampled packet, in bytes
- **Intf num**—Unique number that identifies the sampled logical interface
- **IP frag**—IP fragment number, if applicable
- **TCP flags**—Any TCP flags found in the IP header

To set the timestamp option for the file **my-sample**, enter the following:

```
[edit forwarding-options sampling output file]
```

```
user@host# set filename my-sample files 5 size 2m world-readable stamp;
```

Whenever you toggle the timestamp option, a new header is included in the file. If you set the **stamp** option, the **Time** field is displayed.

```
# Apr  7 15:48:50
# Time           Dest      Src  Dest  Src Proto  TOS   Pkt  Intf   IP   TCP
#                addr      addr port  port                len   num  frag  flags
# Feb  1 20:31:21
#                Dest      Src  Dest  Src Proto  TOS   Pkt  Intf   IP   TCP
#                addr      addr port  port                len   num  frag  flags
```

Tracing Traffic Sampling Operations

Tracing operations track all traffic sampling operations and record them in a log file in the `/var/log` directory. By default, this file is named `/var/log/sampled`. The default file size is 128 KB, and 10 files are created before the first one gets overwritten.

To trace traffic sampling operations, include the **traceoptions** statement at the `[edit forwarding-options sampling]` hierarchy level:

```
[edit forwarding-options sampling]
traceoptions {
  file filename {
    files number;
    size bytes;
    (world-readable | no-world-readable);
  }
}
```

Configuring Flow Aggregation (cflowd)

You can collect an aggregate of sampled flows and send the aggregate to a specified host that runs the `cflowd` application available from the Cooperative Association for Internet Data Analysis (CAIDA) (<http://www.caida.org>). By using `cflowd`, you can obtain various types of byte and packet counts of flows through a router.

The `cflowd` application collects the sampled flows over a period of 1 minute. At the end of the minute, the number of samples to be exported are divided over the period of another minute and are exported over the course of the same minute.

Before you can perform flow aggregation, the routing protocol process must export the autonomous system (AS) path and routing information to the sampling process. To do this, include the **route-record** statement:

```
route-record;
```

You can include this statement at the following hierarchy levels:

- `[edit routing-options]`
- `[edit routing-instances routing-instance-name routing-options]`

By default, flow aggregation is disabled. To enable the collection of flow aggregates, include the `cflowd` statement at the [edit forwarding-options sampling output] hierarchy level:

```
[edit forwarding-options sampling output]
cflowd hostname {
  aggregation {
    autonomous-system;
    destination-prefix;
    protocol-port;
    source-destination-prefix {
      caida-compliant;
    }
    source-prefix;
  }
  autonomous-system-type (origin | peer);
  (local-dump | no-local-dump);
  port port-number;
  source-address address;
  version format;
}
```

In the `cflowd` statement, specify the name, identifier, and source-address of the host that collects the flow aggregates. You must also include the UDP port number on the host and the **version**, which gives the format of the exported `cflowd` aggregates. To specify an IPv4 source address, include the **source-address** statement. To collect `cflowd` records in a log file before exporting, include the **local-dump** statement. To specify the `cflowd` version number, include the **version** statement. The `cflowd` version is either 5 or 8.

You can specify both host (`cflowd`) sampling and port mirroring in the same configuration. You can perform RE-sampling and port mirroring actions simultaneously. However, you cannot perform PIC-sampling and port mirroring actions simultaneously.

To specify aggregation of specific types of traffic, include the **aggregation** statement. This conserves memory and bandwidth enabling `cflowd` to export targeted flows rather than all the aggregated



NOTE: Aggregation is valid only if `cflowd` version 8 is specified.

To specify a flow type, include the **aggregation** statement at the [edit forwarding-options sampling output `cflowd hostname`] hierarchy level:

```
[edit forwarding-options sampling output cflowd hostname]
aggregation {
  source-destination-prefix;
}
```

You specify the aggregation type using one of the following options:

- **autonomous-system**—Aggregate by AS number; may require setting the separate `cflowd autonomous-system-type` statement to include either **origin** or **peer** AS numbers. The **origin** option specifies to use the origin AS of the packet source address in the Source Autonomous System `cflowd` field. The **peer** option specifies to use the peer AS through which the packet passed in the Source Autonomous System `cflowd` field. By default, **cflowd** exports the origin AS number.
- **destination-prefix**—Aggregate by destination prefix (only).
- **protocol-port**—Aggregate by protocol and port number; requires setting the separate `cflowd port` statement.
- **source-destination-prefix**—Aggregate by source and destination prefix. Version 2.1b1 of CAIDA's `cflowd` application does not record source and destination mask length values in compliance with CAIDA's *cflowd Configuration Guide*, dated August 30, 1999. If you configure the **caida-compliant** statement, the JUNOS software complies with Version 2.1b1 of `cflowd`. If you do not include the **caida-compliant** statement in the configuration, the JUNOS software records source and destination mask length values in compliance with the *cflowd Configuration Guide*.
- **source-prefix**—Aggregate by source prefix (only).

Collection of sampled packets in a local ASCII file is not affected by the **cflowd** statement.

Debugging cflowd Flow Aggregation

To collect the `cflowd` flows in a log file before they are exported, include the **local-dump** option at the `[edit forwarding-options sampling output cflowd hostname]` hierarchy level:

```
[edit forwarding-options sampling output cflowd hostname]
local-dump;
```

By default, the flows are collected in `/var/log/sampled`; to change the filename, include the **filename** statement at the `[edit forwarding-options sampling traceoptions]` hierarchy level. For more information about changing the filename, see “Configuring Traffic Sampling Output” on page 324.



NOTE: Because the **local-dump** option adds extra overhead, you should use it only while debugging `cflowd` problems, not during normal operation.

The following is an example of the flow information. The AS number exported is the origin AS number. All flows that belong under a `cflowd` header are dumped, followed by the header itself:

```
Jun 27 18:35:43 v5 flow entry
Jun 27 18:35:43   Src addr: 10.53.127.1
Jun 27 18:35:43   Dst addr: 10.6.255.15
Jun 27 18:35:43   Nhop addr: 192.168.255.240
Jun 27 18:35:43   Input interface: 5
Jun 27 18:35:43   Output interface: 3
Jun 27 18:35:43   Pkts in flow: 15
```

```

Jun 27 18:35:43 Bytes in flow: 600
Jun 27 18:35:43 Start time of flow: 7230
Jun 27 18:35:43 End time of flow: 7271
Jun 27 18:35:43 Src port: 26629
Jun 27 18:35:43 Dst port: 179
Jun 27 18:35:43 TCP flags: 0x10
Jun 27 18:35:43 IP proto num: 6
Jun 27 18:35:43 TOS: 0xc0
Jun 27 18:35:43 Src AS: 7018
Jun 27 18:35:43 Dst AS: 11111
Jun 27 18:35:43 Src netmask len: 16
Jun 27 18:35:43 Dst netmask len: 0

```

[... 41 more v5 flow entries; then the following header:]

```

Jun 27 18:35:43 cflowd header:
Jun 27 18:35:43 Num-records: 42
Jun 27 18:35:43 Version: 5
Jun 27 18:35:43 Flow seq num: 118
Jun 27 18:35:43 Engine id: 0
Jun 27 18:35:43 Engine type: 3

```

Configuring Active Flow Monitoring Using Version 9

Beginning with JUNOS Release 8.3, you can collect a record of sampled flows using the version 9 format as defined in RFC 3954, *Cisco Systems NetFlow Services Export Version 9*. Version 9 uses templates to collect an set of sampled flows and send the record to a specified host.

You configure the version 9 template used to collect a record of sampled flows at the [edit services monitoring] hierarchy level. For more information, see the *JUNOS Services Interfaces Configuration Guide* and the *JUNOS Feature Guide*.

To enable the collection of traffic flows using the version 9 format, include the **version9** statement at the [edit forwarding-options sampling output cflowd *hostname*] hierarchy level:

```

[edit forwarding-options sampling output cflowd hostname]
version9 {
  template template-name;
}

```

template-name is the name of the version 9 template configured at the [edit services monitoring] hierarchy level.

You configure traffic sampling at the [edit forwarding-options sampling input] hierarchy level. Beginning with JUNOS Release 8.3, traffic sampling supports MPLS traffic as well as IPv4 traffic. You can define a version 9 flow record template suitable for IPv4 traffic, MPLS traffic, or a combination of the two. However, you can sample packets from only one type of family (**inet** or **mpls**) at the same time. For more information about how to configure traffic sampling, see “Configuring Traffic Sampling” on page 288.

The following restrictions apply to configuration of the version 9 format:

- You can configure only one host to collect traffic flows using the version 9 format. Configure the host at the [edit forwarding-options sampling output cflowd *hostname*] hierarchy level.
- You cannot specify both the version 9 format and cflowd versions 5 and 8 formats in the same configuration. For more information about how to configure flow monitoring using cflowd version 8, see “Configuring Flow Aggregation (cflowd)” on page 326.
- Any values for `flow-active-timeout` and `flow-inactive-timeout` that you configure at the [edit forwarding-options sampling output] hierarchy level are overridden by the values configured in the version 9 template.
- Version 9 does not support Routing Engine-based sampling. You cannot configure version 9 to send traffic sampling result to a file in the `/var/tmp` directory.

Example: Configuring Active Flow Monitoring Using Version 9

In this example, you enable active flow monitoring using version 9. You specify a template `mpls` that you configure at [edit services monitoring] hierarchy level. You also configure the traffic family `mpls` to sample.

```
[edit forwarding-options]
sampling {
  input {
    family mpls {
      rate 1;
      run-length;
    }
  }
  output {
    cflowd 10.60.2.1 { # The IP address and port of the host
      port 2055; # that collects the sampled traffic flows.
      source-address 3.3.3.1;
      version9 {
        template mpls; # Version 9 records are sent
      } # using the template named mpls
    }
  }
}
```

Configuring Port Mirroring

On routing platforms containing an Internet Processor II ASIC, you can send a copy of an IPv4 packet from the routing platform to an external host address or a packet analyzer for analysis. This is known as *port mirroring*.

Port mirroring is different from traffic sampling. In traffic sampling, a sampling key based on the packet header is sent to the Routing Engine. There, the key can be placed in a file, or cflowd packets based on the key can be sent to a cflowd server. In port mirroring, the entire packet is copied and sent out through a next-hop interface.

One application for port mirroring sends a duplicate packet to a virtual tunnel. Then, a next-hop group can be configured to forward copies of this duplicate packet to

several interfaces. For more information about next-hop groups, see “Configuring a Next-Hop Group” on page 291.

To configure port mirroring, include the **port-mirroring** statement at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
port-mirroring {
  family (inet | inet6) {
    input {
      rate number;
      run-length number;
    }
    output {
      interface interface-name {
        next-hop address;
      }
      no-filter-check;
    }
  }
}
```

To configure port mirroring, include the **port-mirroring** statement. To configure the address family, rate of sampling, and length of sampling for port mirroring, include the **input** statement. To specify on which interface to send duplicate packets and the next-hop address to send packets, include the **output** statement. To see if there are any filters on the specified interface, include the **no-filter-check** statement.

For information about the **rate** and **run-length** statements, see “Configuring Traffic Sampling” on page 288.

In typical applications, you send the sampled packets to an analyzer or a workstation for analysis, not to another router. If you must send this traffic over a network, you should use tunnels. For more information about tunnel interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

Beginning with JUNOS release 9.3, port mirroring is supported for Layer 2 traffic on MX-series routers only. For information about how to configure, see the *MX-series Layer 2 Configuration Guide*.

The following restrictions apply:

- You can configure port mirroring for IPv4 and IPv6 simultaneously on the M120 and M320 routers and the MX-series routers.
- You can configure either IPv4 or IPv6 port mirroring but not both on M-series routers, except for the M120 and M320 routers, which support port mirroring for IPv4 and IPv6 simultaneously.
- You cannot configure firewall filters on the port-mirroring interface.
- The interface you configure for port mirroring should not participate in any kind of routing activity.
- The destination address you specify should not have a route to the ultimate traffic destination. For example, if the sampled IPv4 packets have a destination address of 190.68.9.10 and the port-mirrored traffic is sent to 190.68.20.15 for analysis,

the device associated with the latter address should not know a route to 190.68.9.10. Also, it should not send the sampled packets back to the source address.

- Only transit data is supported.
- You can configure only one port-mirroring interface per router. If you include more than one interface in the **port-mirroring** statement, the previous one is overwritten.
- You must include a firewall filter with both the **accept** action and the **port-mirror** action modifier on the inbound interface. Do not include the **discard** action, or port mirroring will not work.

Configuring Packet Capture

Packet capture allows you to monitor and analyze offline IP version 4 (IPv4) packets flowing through a router. Packet capture monitors packet fragments also. Packet capture can be enabled on any interface and can analyze ingress traffic, egress traffic, or both.



NOTE: Packet capture is supported for the J-series Services Routers only. Packet capture is not supported on tunnel interfaces. You cannot configure packet capture and sampling at the same time.

To configure packet capture, include the **packet-capture** statement at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
packet-capture {
  disable;
  file {
    filename file-name;
    files number;
    size number;
    (world-readable | no-world-readable);
  }
  maximum-capture-size number;
}
```

To disable packet capture, include the **disable** statement. Packet capture is enabled by default.

You can capture packets into files. Files are classified based on the physical interface the packets are captured on (one file per physical interface). You can specify the file name, maximum size, and maximum number of files. When you capture a file named **pcap-file**, packet capture creates one file for each physical interface and appends the physical interface designator to the filename (for example, **at**). When the file named **pcap-file.xx** reaches its maximum size, the file is renamed **pcap-file.xx.0**. When **pcap-file.xx** reaches its maximum size again, the file is renamed **pcap-file.xx.1**. This process continues until the maximum number of files is exceeded. When that happens, the oldest file is overwritten. The file named **pcap-file.xx** is always the latest

file. The packet capture file for an interface is created when the first packet is captured on that interface. Once created, this file is not removed even if packet capture is disabled on the interface. All packet capture files are stored in the `/var/tmp/` directory.

If the PCAP file is deleted from `var/tmp/`, the file is not recreated upon the next packet capture traffic on the interface. You must first disable and then enable PCAP functionality again to recreate the PCAP file.

To enable capture into files, include the `file` statement. You can specify the target filename, maximum file size, and the maximum number of files. To specify the name of the target file, include the `filename` statement. To specify the maximum size of the file, include the `size` statement. To specify the maximum number of files, include the `files` statement.

To specify the maximum size of the packet for capture, include the `maximum-capture-size` statement.

You can capture packets on a specific interface by configuring either of the following:

- Configure a firewall filter with the action `sample` and apply it to the interface.
- Configure sampling on the interface in the ingress or egress traffic.



NOTE: Interface sampling does not capture host-originated packets. Configure firewall filters to capture host-originated packets.



NOTE: Firewall filter applied to a loopback interface (`lo0`) affects all packets going to and from the Routing Engine.

You can capture packets on a specific interface. For information about configuring interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

You can capture only specific types of packets by using a firewall filter in conjunction with packet capture. To configure packet capture for specific packets using firewall filters, include the following statements at the `[edit firewall]` hierarchy level:

```
[edit firewall]
filter filter-name {
  term term-name {
    from {
      match-conditions;
    }
    then {
      sample;
      accept;
    }
  }
}
```



NOTE: Configure packet capture with appropriate firewall filters to control the number of packets captured. Performance of the router may be impacted if packet capture is used without configuring any firewall filters.



NOTE: Packet capture does not support multilink encapsulations (such as MLPPP).

You must disable packet capture to modify encapsulation. To modify the encapsulation on a packet capture-enabled interface, perform the following tasks:

1. Disable packet capture by including the **disable** statement at the [edit forwarding-options packet-capture] hierarchy level.
2. Remove the packet capture file for the interface from the `/var/tmp/` directory.
3. Change the encapsulation.
4. Enable packet capture.

For packets captured on T1, T3, E1, E3, SE, and ISDN interfaces in the egress direction, the size of packets captured can be one byte less than the configured value of **maximum-capture-size** because of the PLP byte.

To capture packets on an ISDN interface, configure packet capture on the dialer interface. To capture packets on the PPPoE interface, configure packet capture on the PPPoE interface.

Packet capture is not supported with MLPPP encapsulation. However, the CLI does not prevent you from enabling packet capture on an interface with MLPPP encapsulation. If packet capture is enabled in the input direction on an interface with MLPPP encapsulation, input packets on that interface are captured on the output interfaces.

By default, there is no tracing operation support for packet capture.

For more information about configuring specific interface types, see the *JUNOS Network Interfaces Configuration Guide*.

Chapter 15

Summary of Traffic Sampling and Forwarding Options Configuration Statements

The following sections explain each of the sampling and forwarding statements. The statements are organized alphabetically.

accounting

Syntax

```

accounting group-name {
  output {
    cflowd [ hostnames ] {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      autonomous-system-type (origin | peer);
      port port-number;
      version format;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface [ interface-names ] {
      engine-id number;
      engine-type number;
      source-address address;
    }
  }
}

```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify discard accounting instance name and options.

The statements are explained separately.

Usage Guidelines See “Configuring Discard Accounting” on page 289.

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

active-server-group

Syntax	<code>active-server-group server-group-name;</code>
Hierarchy Level	<code>[edit forwarding-options dhcp-relay],</code> <code>[edit forwarding-options dhcp-relay group group-name],</code> <code>[edit logical-systems logical-system-name forwarding-options dhcp-relay],</code> <code>[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name</code> <code>forwarding-options dhcp-relay],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name</code> <code>forwarding-options dhcp-relay group group-name],</code> <code>[edit routing-instances routing-instance-name forwarding-options dhcp-relay],</code> <code>[edit routing-instances routing-instance-name forwarding-options dhcp-relay group</code> <code>group-name]</code>
Release Information	Statement introduced in JUNOS Release 8.3.
Description	<p>Apply a DHCP relay agent configuration to the named group of DHCP server addresses.</p> <p>You can include the active-server-group statement at the <code>[edit forwarding-options dhcp-relay]</code> hierarchy level as a global DHCP relay agent configuration option, or at the <code>[edit forwarding-options dhcp-relay group group-name]</code> hierarchy level as a DHCP relay agent configuration option that applies only to a named group of interfaces.</p> <p>Including the active-server-group statement at the <code>[edit forwarding-options dhcp-relay group group-name]</code> hierarchy level as a group-specific option overrides use of the active-server-group statement at the <code>[edit forwarding-options dhcp-relay]</code> hierarchy level as a global option.</p>
Options	<i>server-group-name</i> —Name of the group of DHCP server addresses to which the DHCP relay agent configuration applies.
Usage Guidelines	See “Configuring the Extended DHCP Relay Agent” on page 299.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

aggregation

Syntax	<pre>aggregation { autonomous-system; destination-prefix; protocol-port; source-destination-prefix { caida-compliant; } source-prefix; }</pre>
Hierarchy Level	[edit forwarding-options accounting output cflowd <i>hostname</i>], [edit forwarding-options sampling output cflowd <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For cflowd version 8 only, specify the type of data to be aggregated; cflowd records and sends only those flows that match the specified criteria.
Options	<p>autonomous-system—Aggregate by autonomous system (AS) number.</p> <p>caida-compliant—Record source and destination mask length values in compliance with the Version 2.1b1 release of the cflowd application from the Cooperative Association for Internet Data Analysis (CAIDA). If this statement is not configured, the JUNOS software records source and destination mask length values in compliance with the <i>cflowd Configuration Guide</i>, dated August 30, 1999.</p> <p>destination-prefix—Aggregate by destination prefix.</p> <p>protocol-port—Aggregate by protocol and port number.</p> <p>source-destination-prefix—Aggregate by source and destination prefix.</p> <p>source-prefix—Aggregate by source prefix.</p>
Usage Guidelines	See “Configuring Flow Aggregation (cflowd)” on page 326.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

always-write-giaddr

Syntax	<code>always-write-giaddr;</code>
Hierarchy Level	[edit forwarding-options dhcp-relay overrides], [edit forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Overwrite the gateway IP address (giaddr) of every DHCP packet with the gateway IP address of the DHCP relay agent before forwarding the packet to the DHCP server.
Usage Guidelines	See “Configuring the Extended DHCP Relay Agent” on page 299.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

always-write-option-82

Syntax	always-write-option-82;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay overrides],</p> <p>[edit forwarding-options group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides]</p>
Release Information	Statement introduced in JUNOS Release 8.3.
Description	<p>Override the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server. The use of this option causes the DHCP relay agent to perform one of the following actions, depending on how it is configured:</p> <ul style="list-style-type: none"> ■ If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server. ■ If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.
Usage Guidelines	See “Configuring the Extended DHCP Relay Agent” on page 299.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

authentication

Syntax	<pre> authentication { password <i>password-string</i>; username-include { circuit-type; delimiter <i>delimiter-character</i>; domain-name <i>domain-name-string</i>; logical-system-name; mac-address; option-60; option-82 [circuit-id] [remote-id]; routing-instance-name; user-prefix <i>user-prefix-string</i>; } }</pre>
Hierarchy Level	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>]</pre>
Release Information	Statement introduced in JUNOS Release 9.1.
Description	<p>Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay or DHCP local server configuration.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Using External AAA Authentication Services” on page 313.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

autonomous-system-type

Syntax	autonomous-system-type (origin peer);
Hierarchy Level	[edit forwarding-options accounting output cflowd <i>hostname</i>], [edit forwarding-options sampling output cflowd <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the type of AS numbers that cflowd exports.
Options	<p>origin—Export origin AS numbers of the packet source address in the Source Autonomous System cflowd field.</p> <p>peer—Export peer AS numbers through which the packet passed in the Source Autonomous System cflowd field.</p> <p>Default: origin</p>
Usage Guidelines	See “Configuring Flow Aggregation (cflowd)” on page 326.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

bootp

Syntax bootp {
 client-response-ttl *number*;
 description *description-of-service*;
 interface *interface-group* {
 client-response-ttl *number*;
 description *description-of-interface*;
 maximum-hop-count *number*;
 minimum-wait-time *seconds*;
 no-listen;
 server [*addresses*];
 }
 maximum-hop-count *number*;
 minimum-wait-time *seconds*;
 server [*addresses*];
 }

Hierarchy Level [edit forwarding-options helpers]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configures a router or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent.

DHCP relaying is disabled.

Options The remaining statements are explained separately.

Usage Guidelines See “Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent” on page 294.

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

cflowd

See the following sections:

- cflowd (Discard Accounting) on page 344
- cflowd (Flow Monitoring) on page 345
- cflowd (Sampling) on page 346

cflowd (Discard Accounting)

Syntax `cflowd hostname {
 aggregation {
 autonomous-system;
 destination-prefix;
 protocol-port;
 source-destination-prefix {
 caida-compliant;
 }
 source-prefix;
 }
 autonomous-system-type (origin | peer);
 port port-number;
 source-address address;
 version format;
}`

Hierarchy Level [edit forwarding-options accounting *group-name* output]

Release Information Statement introduced before JUNOS Release 7.4.

Description Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility `cfcollect`.

You can configure up to one version 5 and one version 8 flow format at the [edit forwarding-options accounting *group-name* output] hierarchy level.

Options *hostname*—The IP address or identifier of the host system (the workstation running the `cflowd` utility).

The remaining statements are explained separately.

Usage Guidelines See “Configuring Flow Aggregation (cflowd)” on page 326.

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

cflowd (Flow Monitoring)

Syntax	<code>cflowd hostname { port <i>port-number</i>; }</code>
Hierarchy Level	[edit forwarding-options monitoring <i>group-name</i> family inet output]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility <code>cfcollect</code>.</p> <p>You can configure up to eight version 5 flow formats at the [edit forwarding-options monitoring <i>group-name</i> output] hierarchy level. Version 8 flow formats are not supported for flow-monitoring applications.</p>
Options	<p><i>hostname</i>—The IP address or identifier of the host system (the workstation running the <code>cflowd</code> utility).</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Flow Monitoring” on page 290.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

cflowd (Sampling)

Syntax

```
cflowd hostname {
  version9 {
    template template-name;
  }
  aggregation {
    autonomous-system;
    destination-prefix;
    protocol-port;
    source-destination-prefix {
      caida-compliant;
    }
    source-prefix;
  }
  autonomous-system-type (origin | peer);
  (local-dump | no-local-dump);
  port port-number;
  source-address address;
  version format;
}
```

Hierarchy Level [edit forwarding-options sampling output]

Release Information Statement introduced before JUNOS Release 7.4.
version9 statement introduced in JUNOS Release 8.3.

Description Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect. Specify a host system to collect sampled flows using the version 9 format.

You can configure up to one version 5 and one version 8 flow format at the [edit forwarding-options sampling output cflowd *hostname*] hierarchy level. For the same configuration, you can specify only either version 9 flow record formats or formats using versions 5 and 8, not both types of formats.

Options *hostname*—The IP address or identifier of the host system (the workstation either running the cflowd utility or collecting traffic flows using version 9.)

You can configure only one host system for version 9.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Flow Aggregation (cflowd)” on page 326 and “Configuring Active Flow Monitoring Using Version 9” on page 329.

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

circuit-id

Syntax	<pre>circuit-id { prefix host-name logical-system-name routing-instance-name; }</pre>
Hierarchy Level	<pre>[edit forwarding-options dhcp-relay relay-option-82], [edit forwarding-options dhcp-relay group group-name relay-option-82], [edit logical-systems logical-system-name forwarding-options dhcp-relay relay-option-82], [edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name relay-option-82], [edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82], [edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay group group-name relay-option-82], [edit routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82], [edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name relay-option-82]</pre>
Release Information	Statement introduced in JUNOS Release 8.3.
Description	<p>Include the agent-circuit-id suboption (suboption 1) of the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server.</p> <p>The format of the agent-circuit-id information for Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual local area networks (VLANs) or stacked VLANs (S-VLANs) is as follows:</p> <pre>(fe ge)-fpc/pic/port</pre> <p>The format of the agent-circuit-id information for Fast Ethernet or Gigabit Ethernet interfaces that use VLANs is as follows:</p> <pre>(fe ge)-fpc/pic/port:vlan-id</pre> <p>The format of the agent-circuit-id information for Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs is as follows:</p> <pre>(fe ge)-fpc/pic/port:svlan-id-vlan-id</pre> <p>The remaining statement is explained separately.</p>
Usage Guidelines	See “Configuring the Extended DHCP Relay Agent” on page 299.
Required Privilege Level	<pre>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</pre>

circuit-type

Syntax	circuit-type;
Hierarchy Level	[edit forwarding-options dhcp-relay authentication username-include], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Specify that the circuit type be concatenated with the username during the subscriber authentication process.
Usage Guidelines	See “Using External AAA Authentication Services” on page 313.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

client-response-ttl

Syntax	client-response-ttl <i>number</i> ;
Hierarchy Level	[edit forwarding-options helpers bootp], [edit forwarding-options helpers bootp interface <i>interface-group</i>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Perform an IP time-to-live (TTL) decrement before accepting any packets for processing when DHCP response packets are sent to a DHCP client.
Options	<i>number</i> —Decrement amount. Default: None
Usage Guidelines	See “Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent” on page 294.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

default-local-server-group

Syntax	<code>default-local-server-group <i>local-server-group-name</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay relay-option-60 vendor-option],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option]</p>
Release Information	Statement introduced in JUNOS Release 9.0.
Description	<p>Forward DHCP client packets to a default extended DHCP local server when you use the DHCP vendor class identifier option (option 60) in DHCP packets to forward client traffic to specific DHCP servers.</p> <p>If the option 60 string received in the DHCP client packet does not match the ASCII or hexadecimal match string and match criteria (exact match or partial match) that you specify, the extended DHCP relay agent forwards the client packets to the specified default DHCP local server group configured with the <code>dhcp-local-server</code> statement at the [edit system services] hierarchy level.</p>
Options	<i>local-server-group-name</i> —Name of the default extended DHCP local server group.
Usage Guidelines	See “Using Option 60 Information to Forward Client Traffic to Specific DHCP Servers” on page 306.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

default-relay-server-group

Syntax	<code>default-relay-server-group <i>server-group-name</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay relay-option-60 vendor-option],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option]</p>
Release Information	Statement introduced in JUNOS Release 9.0.
Description	<p>Relay DHCP client packets to a default group of extended DHCP relay servers when you use the DHCP vendor class identifier option (option 60) in DHCP packets to forward client traffic to specific DHCP servers.</p> <p>If the option 60 string received in the DHCP client packet does not match the ASCII or hexadecimal match string and match criteria (exact match or partial match) that you specify, the extended DHCP relay agent relays the client packets to the specified default group of servers configured with the server-group statement at the [edit forwarding-options dhcp-relay] hierarchy level.</p>
Options	<i>server-group-name</i> —Name of the default DHCP relay server group.
Usage Guidelines	See “Using Option 60 Information to Forward Client Traffic to Specific DHCP Servers” on page 306.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

delimiter

Syntax	<code>delimiter <i>delimiter-character</i>;</code>
Hierarchy Level	[edit forwarding-options dhcp-relay authentication username-include], [edit forwarding-options dhcp-relay group authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Specify the character used as the delimiter between the concatenated components of the username. You cannot use the semicolon (;) as a delimiter.
Usage Guidelines	See “Using External AAA Authentication Services” on page 313.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

description

See the following sections:

- [description \(Interface\)](#) on page 352
- [description \(Service\)](#) on page 352

description (Interface)

Syntax	description <i>description-of-interface</i> ;
Hierarchy Level	[edit forwarding-options helpers bootp interface <i>interface-group</i>], [edit forwarding-options helpers domain interface <i>interface-name</i>], [edit forwarding-options helpers tftp interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Text description of interface.
Usage Guidelines	See “Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent” on page 294 and “Configuring DNS and TFTP Packet Forwarding” on page 295.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

description (Service)

Syntax	description <i>description-of-service</i> ;
Hierarchy Level	[edit forwarding-options helpers bootp], [edit forwarding-options helpers domain], [edit forwarding-options helpers tftp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Description of BOOTP, DHCP, Domain Name System (DNS), or Trivial File Transfer Protocol (TFTP) service.
Usage Guidelines	See “Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent” on page 294 and “Configuring DNS and TFTP Packet Forwarding” on page 295.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dhcp-relay

```

Syntax  dhcp-relay {
            authentication {
                password password-string;
                username-include {
                    circuit-type;
                    delimiter delimiter-character;
                    domain-name domain-name-string;
                    logical-system-name;
                    mac-address;
                    option-60;
                    option-82 [circuit-id] [remote-id];
                    routing-instance-name;
                    user-prefix user-prefix-string;
                }
            }
            overrides {
                always-write-giaddr;
                always-write-option-82;
                layer2-unicast-replies;
                trust-option-82;
                disable-relay;
            }
            relay-option-60 {
                vendor-option {
                    (equals | starts-with) (ascii match-string | hexadecimal match-hex) {
                        (relay-server-group server-group-name | local-server-group local-server-group-name
                        | drop);
                    }
                    (default-relay-server-group server-group-name | default-local-server-group
                    local-server-group-name | drop);
                }
            }
            relay-option-82 {
                circuit-id {
                    prefix host-name logical-system-name routing-instance-name;
                }
            }
            server-group {
                server-group-name {
                    server-ip-address;
                }
            }
            active-server-group server-group-name;
            group group-name {
                active-server-group server-group-name;
                authentication {
                    password password-string;
                    username-include {
                        circuit-type;
                        delimiter delimiter-character;
                        domain-name domain-name-string;

```

```

        logical-system-name;
        mac-address;
        option-60;
        option-82 [circuit-id] [remote-id];
        routing-instance-name;
        user-prefix user-prefix-string;
    }
}
overrides {
    always-write-giaddr;
    always-write-option-82;
    layer2-unicast-replies;
    trust-option-82;
    disable-relay;
}
relay-option-60 {
    vendor-option {
        (equals | starts-with) (ascii match-string | hexadecimal match-hex) {
            (relay-server-group server-group-name |
             local-server-group local-server-group-name |
             drop);
        }
        (default-relay-server-group server-group-name |
         default-local-server-group local-server-group-name |
         drop);
    }
}
relay-option-82 {
    circuit-id {
        prefix host-name logical-system-name routing-instance-name;
    }
}
interface interface-name [upto upto-interface-name] [exclude];
}
traceoptions {
    flag all;
    flag database;
    flag state;
    flag interface;
    flag rtsock;
    flag packet;
    flag packet-option;
    flag io;
    flag ha;
    flag ui;
    flag general;
    flag fwd;
    flag rpd;
    file file-name {
        <files number>;
        <size maximum-file-size>;
        <match regex>;
        <world-readable | no-world-readable>;
    }
}
}

```


Hierarchy Level	[edit forwarding-options], [edit logical-systems <i>logical-system-name</i> forwarding-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options], [edit routing-instances <i>routing-instance-name</i> forwarding-options]
Release Information	Statement introduced in JUNOS Release 8.3. traceoptions option introduced in JUNOS Release 8.5. relay-option-60 option introduced in JUNOS Release 9.0. authentication option introduced in JUNOS Release 9.1.
Description	<p>Configure extended Dynamic Host Configuration Protocol (DHCP) relay options on the router and enable the router to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.</p> <p>The extended DHCP relay agent options configured with the <code>dhcp-relay</code> statement are incompatible with the DHCP/BOOTP relay agent options configured with the <code>bootp</code> statement. As a result, the extended DHCP relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router at the same time.</p> <p>The extended DHCP relay interacts with the local AAA service framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication. You can configure authentication support on a global basis or for a specific group of interfaces.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring the Extended DHCP Relay Agent” on page 299 and “Using External AAA Authentication Services” on page 313.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

disable

See the following sections:

- **disable (Packet Capture)** on page 356
- **disable (Sampling)** on page 356

disable (Packet Capture)

Syntax	disable;
Hierarchy Level	[edit forwarding-options packet-capture],
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Disable packet capture.
Usage Guidelines	See “Configuring Packet Capture” on page 332.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

disable (Sampling)

Syntax	disable;
Hierarchy Level	[edit forwarding-options sampling], [edit forwarding-options sampling output file]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable traffic sampling.
Usage Guidelines	See “Disabling Traffic Sampling” on page 321.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

disable-relay

Syntax	disable-relay;
Hierarchy Level	[edit forwarding-options dhcp-relay overrides], [edit forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Disable DHCP relay on specific interfaces in a group.
Usage Guidelines	See “Configuring the Extended DHCP Relay Agent” on page 299.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

domain

Syntax	domain { description <i>description-of-service</i> ; server address < [routing-instance <i>routing-instance-names</i>] >; interface <i>interface-name</i> { description <i>description-of-interface</i> ; no-listen; server address < [routing-instance <i>routing-instance-names</i>] >; } }
Hierarchy Level	[edit forwarding-options helpers]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable DNS request packet forwarding.
Options	The statements are explained separately.
Usage Guidelines	See “Configuring DNS and TFTP Packet Forwarding” on page 295.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

domain-name

Syntax	<code>domain-name <i>domain-name-string</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay authentication username-include], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]</p>
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Specify the domain name that is concatenated with the username during the subscriber authentication process.
Options	<i>domain-name-string</i> —The domain name formatted string.
Usage Guidelines	See “Using External AAA Authentication Services” on page 313.
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>

drop

Syntax	drop;
Hierarchy Level	<pre> [edit forwarding-options dhcp-relay relay-option-60 vendor-option], [edit forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)], [edit forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option], [edit forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)]</pre>
Release Information	Statement introduced in JUNOS Release 9.0.
Description	<p>Drop (discard) DHCP client packets when you use the DHCP vendor class identifier option (option 60) in DHCP packets to forward client traffic to specific DHCP servers.</p> <p>To drop DHCP client packets that contain an option 60 string that matches the ASCII or hexadecimal match string and match criteria (exact match or partial match) that you specify, include the drop statement at the [edit forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)] hierarchy level.</p> <p>To drop DHCP client packets that contain an option 60 string that does <i>not</i> match the ASCII or hexadecimal match string and match criteria (exact match or partial</p>

match) that you specify, include the **drop** statement at the [edit forwarding-options dhcp-relay relay-option-60 vendor-option] hierarchy level.

Usage Guidelines See “Using Option 60 Information to Forward Client Traffic to Specific DHCP Servers” on page 306.

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

export-format

Syntax export-format cflowd-version-5;

Hierarchy Level [edit forwarding-options monitoring *group-name* family inet output]

Release Information Statement introduced before JUNOS Release 7.4.

Description Flow monitoring export format.

Options cflowd-version-5—Cflowd version 5.

Usage Guidelines See “Configuring Flow Monitoring” on page 290.

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

family

See the following sections:

- family (Filtering) on page 361
- family (Sampling) on page 362

family (Filtering)

Syntax family *family-name* {
 filter {
 input *filter-name*;
 output *filter-name*;
 }
 flood {
 input *filter-name*;
 }
 route-accounting;
 }

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before JUNOS Release 7.4.
 route-accounting option introduced in JUNOS Release 8.3; supported only with IPv6.

Description Specify address family for filters.

Options *family-name*—Address family. Specify **inet** for IP version 4 (IPv4), **inet6** for IP version 6 (IPv6), **mpls** for Multiprotocol Label Switching (MPLS), or **vppls** for virtual private LAN service (VPLS).

The remaining statements are explained separately.

Usage Guidelines See “Configuring a Forwarding Table Filter” on page 286.

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

family (Sampling)

Syntax	family (inet mpls) { max-packets-per-second <i>number</i> ; rate <i>number</i> ; run-length <i>number</i> ; }
Hierarchy Level	[edit forwarding-options sampling input]
Release Information	Statement introduced before JUNOS Release 7.4. mpls option introduced in JUNOS Release 8.3.
Description	Configure the protocol family to be sampled.
Options	inet—IP version 4 (IPv4) mpls—Multiprotocol Label Switching (MPLS) The remaining statements are explained separately.
Usage Guidelines	See “Configuring Traffic Sampling” on page 288.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

family inet

See the following sections:

- family inet (Load Balancing) on page 363
- family inet (Monitoring) on page 364

family inet (Load Balancing)

Syntax family inet {
 layer-3;
 layer-4;
}

Hierarchy Level [edit forwarding-options hash-key]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure layer information for the load-balancing specification. Only the IPv4 protocol is supported.

Options layer-3—Include Layer 3 (IP) data in the hash key.

layer-4—Include Layer 4 Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) data in the hash key.

Usage Guidelines See “Configuring the Load-Balance Per-Packet Action” on page 136.

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

family inet (Monitoring)

Syntax

```
family inet {
  output {
    cflowd hostname {
      port port-number;
    }
    export-format cflowd-version-5;
    flow-active-timeout seconds;
    flow-export-destination {
      (cflowd-collector | collector-pic);
    }
    flow-inactive-timeout seconds;
    interface interface-name {
      engine-id number;
      engine-type number;
      input-interface-index number;
      output-interface-index number;
      source-address address;
    }
  }
}
```

Hierarchy Level [edit forwarding-options monitoring *group-name*]

Release Information Statement introduced before JUNOS Release 7.4

Description Configure flow monitoring for an address family. Only the IPv4 protocol is supported.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Flow Monitoring” on page 290

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

family mpls

Syntax

```
family mpls {
  label-1;
  label-2;
  label-3;
  no-labels;
  no-label-1-exp;
  payload {
    ip {
      layer-3-only;
      port-data {
        source-msb;
        source-lsb;
        destination-msb;
        destination-lsb;
      }
    }
  }
}
```

Hierarchy Level [edit forwarding-options hash-key]

Release Information Statement introduced before JUNOS Release 7.4.
 no-label-1-exp option introduced in JUNOS Release 8.0.
 label-3 and no-labels options introduced in JUNOS Release 8.1.

Description For aggregated Ethernet and SONET/SDH interfaces only, configure load balancing based on MPLS labels. Only the IPv4 protocol is supported.

Options

- label-1—Include only one label in the hash key.
- label-2—Include both labels in the hash key.
- label-3—Include the third MPLS label in the hash key.
- no-labels—Include no MPLS labels in the hash key.
- no-label-1-exp—Do not use the EXP bit of the first label in the hash calculation.
- payload—Include bits from IP payload in the hash key.
- ip—Include the IP address of the IPv4 or IPv6 payload in the hash key.
- layer-3-only—Include only Layer 3 IP information.
- port-data—Include the source and destination port field information.
- source-msb—Include the most significant byte of the source port.
- source-lsb—Include the least significant byte of the source port.

destination-msb—Include the most significant byte of the destination port.

destination-lsb—Include the least significant byte of the destination port.

Usage Guidelines See “Configuring the Load-Balance Per-Packet Action” on page 136.

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

family multiservice

Syntax family multiservice {
 destination-mac;
 source-mac;
}

Hierarchy Level [edit forwarding-options hash-key]

Release Information Statement introduced in JUNOS Release 8.0.

Description (M-series, MX-series, and T-series routing platforms only) Configure load balancing based on Layer 2 media access control information.

Options **destination-mac**—Include the destination-address MAC information in the hash key.
source-mac—Include the source-address MAC information in the hash key.

Usage Guidelines See “Load Balancing Based on Layer 2 MAC Information” on page 140.

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

file

See the following sections:

- file (Helpers Trace Options) on page 367
- file (Packet Capture) on page 368
- file (Sampling) on page 368
- file (Trace Options) on page 369

file (Helpers Trace Options)

Syntax file *filename* {
 files *number*;
 size (Sampling) *bytes*;
 }

Hierarchy Level [edit forwarding-options helpers traceoptions]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure information about the DNS and TFTP packet-forwarding files that contain trace logging information.

Options *filename*—Name of the file containing the trace information.
 Default: /var/log/sampled

The remaining statements are explained separately.

Usage Guidelines See “Tracing BOOTP, DNS, and TFTP Forwarding Operations” on page 296.

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

file (Packet Capture)

Syntax file {
 filename *file-name*;
 files *number*;
 size *number*;
 }

Hierarchy Level [edit forwarding-options packet-capture]

Release Information Statement introduced in JUNOS Release 7.5.

Description Enable packet capture to a file.

Options The statements are explained separately.

Usage Guidelines See “Configuring Packet Capture” on page 332.

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

file (Sampling)

Syntax file {
 disable;
 filename *filename*;
 files *number*;
 size (Sampling) *bytes*;
 (stamp | no-stamp);
 (world-readable | no-world-readable);
 }

Hierarchy Level [edit forwarding-options sampling output]

Release Information Statement introduced before JUNOS Release 7.4.

Description Collect the traffic samples in a file.

The statements are explained separately.

Usage Guidelines See “Configuring Traffic Sampling Output” on page 324.

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

file (Trace Options)

Syntax	file <i>filename</i> { files <i>number</i> ; size (Sampling) <i>bytes</i> ; (world-readable no-world-readable); }
Hierarchy Level	[edit forwarding-options port-mirroring traceoptions], [edit forwarding-options sampling traceoptions]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure information about the files that contain trace logging information.
Options	<i>filename</i> —The name of the file containing the trace information. Default: /var/log/sampled The remaining statements are explained separately.
Usage Guidelines	See “Tracing Traffic Sampling Operations” on page 326.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

filename

See the following sections:

- filename (Packet Capture) on page 370
- filename (Sampling) on page 370

filename (Packet Capture)

Syntax	filename <i>filename</i> ;
Hierarchy Level	[edit forwarding-options packet-capture file]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Configure the name of the output file.
Options	<i>filename</i> —Name of the file.
Usage Guidelines	See “Configuring Packet Capture” on page 332.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

filename (Sampling)

Syntax	filename <i>filename</i> ;
Hierarchy Level	[edit forwarding-options sampling output file]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the name of the output file.
Options	<i>filename</i> —Name of the file in which to place the traffic samples. All files are placed in the directory /var/tmp.
Usage Guidelines	See “Configuring Traffic Sampling Output” on page 324.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

files

See the following sections:

- files (Packet Capture) on page 371
- files (Sampling) on page 371

files (Packet Capture)

Syntax	<code>files number;</code>
Hierarchy Level	[edit forwarding-options packet-capture file]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Configure the maximum number of files for packet capturing.
Options	<i>number</i> —Maximum number of files. Range: 2 through 10,000 files Default: 10
Usage Guidelines	See “Configuring Packet Capture” on page 332.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

files (Sampling)

Syntax	<code>files number;</code>
Hierarchy Level	[edit forwarding-options port-mirroring traceoptions file], [edit forwarding-options sampling output file], [edit forwarding-options sampling traceoptions file]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the total number of files to be saved with samples or trace data.
Options	<i>number</i> —Maximum number of traffic sampling or trace log files. When a file named <i>sampling-file</i> reaches its maximum size, it is renamed <i>sampling-file.0</i> , then <i>sampling-file.1</i> , and so on, until the maximum number of traffic sampling files is reached. Then the oldest sampling file is overwritten. Range: 1 through 100 files Default: 5 files for sampling output; 10 files for trace log information
Usage Guidelines	See “Configuring Traffic Sampling Output” on page 324 and “Tracing Traffic Sampling Operations” on page 326.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

filter

Syntax	filter input <i>filter-name</i> ;
Hierarchy Level	[edit forwarding-options family <i>family-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Apply a forwarding table filter to a forwarding table.
Options	input <i>filter-name</i> —Name of the forwarding table filter.
Usage Guidelines	See “Configuring a Forwarding Table Filter” on page 286.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

flood

Syntax	flood input <i>filter-name</i> ;
Hierarchy Level	[edit forwarding-options family <i>family-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Apply a forwarding table filter to a flood table.
Options	input <i>filter-name</i> —Name of the forwarding table filter.
Usage Guidelines	See “Configuring a Forwarding Table Filter” on page 286.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

flow-active-timeout

Syntax	flow-active-timeout <i>seconds</i> ;
Hierarchy Level	[edit forwarding-options accounting <i>group-name</i> output], [edit forwarding-options monitoring <i>group-name</i> family inet output], [edit forwarding-options sampling output]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the interval before exporting an active flow.
Options	<i>seconds</i> —Interval, in seconds. Range: 60 through 1800 Default: 1800
Usage Guidelines	See “Configuring Discard Accounting” on page 289, “Configuring Flow Monitoring” on page 290, and “Configuring Traffic Sampling Output” on page 324.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

flow-export-destination

Syntax	flow-export-destination { (cflowd-collector collector-pic); }
Hierarchy Level	[edit forwarding-options monitoring <i>group-name</i> family inet output]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure flow collection.
Options	cflowd-collector—cflowd collector. collector-pic—Collector PIC.
Usage Guidelines	See “Configuring Flow Monitoring” on page 290.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

flow-inactive-timeout

Syntax	flow-inactive-timeout <i>seconds</i> ;
Hierarchy Level	[edit forwarding-options accounting <i>group-name</i> output], [edit forwarding-options monitoring <i>group-name</i> family inet output], [edit forwarding-options sampling output]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the interval before a flow is considered inactive.
Options	<i>seconds</i> —Interval, in seconds. Range: 15 through 1800 Default: 60
Usage Guidelines	See “Configuring Discard Accounting” on page 289, “Configuring Flow Monitoring” on page 290, and “Configuring Traffic Sampling Output” on page 324.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

forwarding-options

Syntax	forwarding-options { ... }
Hierarchy Level	[edit]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure traffic forwarding. The statements are explained separately.
Usage Guidelines	See “Traffic Sampling and Forwarding Configuration” on page 279.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

group

Syntax

```

group group-name {
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      logical-system-name;
      mac-address;
      option-60;
      option-82 [ circuit-id ] [ remote-id ];
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  interface interface-name [upto upto-interface-name] [exclude];
  overrides {
    always-write-giaddr;
    always-write-option-82;
    layer2-unicast-replies;
    trust-option-82;
    disable-relay;
  }
  relay-option-60 {
    vendor-option {
      (equals | starts-with) (ascii match-string | hexadecimal match-hex) {
        (relay-server-group server-group-name |
         local-server-group local-server-group-name |
         drop);
      }
      (default-relay-server-group server-group-name |
       default-local-server-group local-server-group-name |
       drop);
    }
  }
  relay-option-82 {
    circuit-id {
      prefix host-name logical-system-name routing-instance-name;
    }
  }
}

```

Hierarchy Level [edit forwarding-options dhcp-relay],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay]

Release Information Statement introduced in JUNOS Release 8.3.

`relay-option-60` option introduced in JUNOS Release 9.0.

Description Specify the name of a group of interfaces that have a common DHCP relay agent configuration. A group must contain at least one interface.

The statements configured at the `[edit forwarding-options dhcp-relay group group-name]` hierarchy level apply only to the named group of interfaces, and override any global DHCP relay agent settings configured with the same statements at the `[edit forwarding-options dhcp-relay]` hierarchy level.

Options *group-name*—Name of a group of interfaces that have a common DHCP relay agent configuration.

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Extended DHCP Relay Agent” on page 299 and “Using External AAA Authentication Services” on page 313.

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

hash-key

Syntax

```

hash-key {
  family inet {
    layer-3;
    layer-4;
  }
  family mpls {
    label-1;
    label-2;
    label-3;
    no-labels;
    no-label-1-exp;
    payload {
      ip {
        layer-3-only;
        port-data {
          destination-lsb;
          destination-msb;
          source-lsb;
          source-msb;
        }
      }
    }
  }
  family multiservice {
    destination-mac;
    source-mac;
  }
}

```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before JUNOS Release 7.4.
no-label-1-exp option introduced in JUNOS Release 8.0.
family multiservice statement introduced in JUNOS Release 8.0.
label-3 and no-labels options introduced in JUNOS Release 8.1.

Description Select which packet header data to use for per-flow load balancing.

Options

- inet—IP address family.
- mpls—MPLS address family.
- multiservice—Multiservice protocol family
- layer-3—Incorporate Layer 3 data into the hash key.
- layer-4—Incorporate Layer 4 data into the hash key.
- no-label-1-exp—The EXP bit of the first label is not used in the hash calculation.

label-1—Incorporate the first label into the hash key.

label-2—Incorporate the second label into the hash key.

label-3—Include the third MPLS label in the hash key.

no-labels—Include no MPLS labels in the hash key.

payload—Incorporate payload data into the hash key.

ip—Include the IP address of the IPv4 or IPv6 payload in the hash key.

layer-3-only—Include only Layer 3 IP information.

port-data—Include the source and destination port field information.

source-msb—Include the most significant byte of the source port.

source-lsb—Include the least significant byte of the source port.

destination-msb—Include the most significant byte of the destination port.

destination-lsb—Include the least significant byte of the destination port.

destination-mac—Include the destination MAC address in the hash key.

Usage Guidelines See “Configuring the Load-Balance Per-Packet Action” on page 136.

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

helpers

Syntax

```

helpers {
  bootp {
    client-response-ttl number;
    description description-of-service;
    interface interface-group {
      client-response-ttl number;
      description description-of-interface;
      maximum-hop-count number;
      minimum-wait-time seconds;
      no-listen;
      server [ addresses ];
    }
    maximum-hop-count number;
    minimum-wait-time seconds;
    server address [ addresses ];
  }
  domain {
    description description-of-service;
    server address < [ routing-instance routing-instance-names ] >;
    interface interface-name {
      description description-of-interface;
      no-listen;
      server address < [ routing-instance routing-instance-names ] >;
    }
  }
  tftp {
    description description-of-service;
    server address < [ routing-instance routing-instance-names ] >;
    interface interface-name {
      description description-of-interface;
      no-listen;
      server address < [ routing-instance routing-instance-names ] >;
    }
  }
  traceoptions {
    file filename {
      files number;
      size bytes;
    }
    level level;
    flag flag;
  }
}

```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description	Enable TFTP or DNS request packet forwarding, or configure the router or interface to act as a DHCP/BOOTP relay agent. Use only one server address per interface or global configuration.
Options	The statements are explained separately.
Usage Guidelines	See “Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent” on page 294 and “Configuring DNS and TFTP Packet Forwarding” on page 295.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

indexed-next-hop

Syntax	indexed-next-hop;
Hierarchy Level	[edit forwarding-options load-balance], [edit logical-systems <i>logical-system-name</i> forwarding-options load-balance], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options load-balance], [edit routing-instances <i>routing-instance-name</i> forwarding-options load-balance]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Generate a permuted index of next-hop entries for unicast and aggregate next hops.
Usage Guidelines	See “Configuring Per-Prefix Load Balancing” on page 293
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

input

See the following sections:

- input (Forwarding Table) on page 381
- input (Port Mirroring) on page 381
- input (Sampling) on page 382

input (Forwarding Table)

Syntax	input <i>filter-name</i> ;
Hierarchy Level	[edit forwarding-options family (inet inet6 mpls) filter], [edit routing-instance <i>instance-name</i> forwarding-options family (inet inet6 mpls) filter]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure filtering on the ingress of the forwarding table.
Options	<i>filter-name</i> —Name of the applied filter.
Usage Guidelines	See “Configuring a Forwarding Table Filter” on page 286.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

input (Port Mirroring)

Syntax	input { <i>rate number</i> ; <i>run-length number</i> ; }
Hierarchy Level	[edit forwarding-options port-mirroring family (inet inet6)]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure port mirroring on a logical interface. The statements are explained separately.
Usage Guidelines	See “Configuring Port Mirroring” on page 330.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

input (Sampling)

Syntax input {
 family (inet | mpls) {
 max-packets-per-second *number* ;
 rate *number*;
 run-length *number*;
 }
 }

Hierarchy Level [edit forwarding-options sampling]

Release Information Statement introduced before JUNOS Release 7.4.
 mpls option introduced in JUNOS Release 8.3.

Description Configure traffic sampling on a logical interface.

The statements are explained separately.

Usage Guidelines See “Configuring Traffic Sampling” on page 288.

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

interface

See the following sections:

- interface (Accounting or Sampling) on page 383
- interface (BOOTP) on page 384
- interface (DHCP Relay Agent) on page 385
- interface (DNS and TFTP Packet Forwarding or Relay Agent) on page 386
- interface (Monitoring) on page 387
- interface (Next-Hop Group) on page 387
- interface (Port Mirroring) on page 388

interface (Accounting or Sampling)

Syntax interface [*interface-names*] {
 engine-id *number*;
 engine-type *number*;
 source-address *address*;
 }

Hierarchy Level [edit forwarding-options accounting *group-name* output],
 [edit forwarding-options sampling output]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the output interface for sending copies of packets elsewhere to be analyzed.

Options *interface-names*—Name of one or more accounting interfaces.
 engine-id number—Identity of the accounting interface.
 engine-type number—Type of this accounting interface.
 source-address address—Address used for generating packets.

Usage Guidelines See “Configuring Discard Accounting” on page 289 and “Configuring Traffic Sampling Output” on page 324.

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

interface (BOOTP)

Syntax interface *interface-group* {
 client-response-ttl *number*;
 description *description-of-interface*;
 maximum-hop-count *number*;
 minimum-wait-time *seconds*;
 no-listen;
 server [*addresses*];
 }

Hierarchy Level [edit forwarding-options helpers bootp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the interface for a DHCP or BOOTP relay agent.

Options *interface-group*—Sets a logical interface or group of logical interfaces with a specific DHCP relay configuration.

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent” on page 294.

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

interface (DHCP Relay Agent)

Syntax	<code>interface <i>interface-name</i> [upto <i>upto-interface-name</i>] [exclude];</code>
Hierarchy Level	[edit forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>]
Release Information	Statement introduced in JUNOS Release 8.3. upto and exclude options introduced in JUNOS Release 9.1.
Description	Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP local server is enabled. You can repeat the interface <i>interface-name</i> statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group. Also, you cannot use an interface that is being used by the DHCP relay agent.
Options	<p>exclude—Exclude an interface or a range of interfaces from the group.</p> <p><i>inteface-name</i>—The name of the interface. You can repeat this keyword multiple times.</p> <p><i>upto-interface-name</i>—The upper end of the range of interfaces; the lower end of the range is the interface-name entry. The interface device name of the <i>upto-interface-name</i> must be the same as the device name of the <i>interface-name</i>.</p>
Usage Guidelines	See “Configuring the Extended DHCP Relay Agent” on page 299 and “Using External AAA Authentication Services” on page 313.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

interface (DNS and TFTP Packet Forwarding or Relay Agent)

Syntax	<pre> interface <i>interface-name</i> { description <i>description-of-service</i>; no-listen; server <i>address</i> < [routing-instance <i>routing-instance-names</i>] >; } </pre>
Hierarchy Level	[edit forwarding-options helpers domain], [edit forwarding-options helpers tftp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the interface for monitoring and forwarding DNS or TFTP requests.
Options	<p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring DNS and TFTP Packet Forwarding” on page 295.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

interface (Monitoring)

Syntax interface *interface-name* {
 engine-id *number*;
 engine-type *number*;
 input-interface-index *number*;
 output-interface-index *number*;
 source-address *address*;
 }

Hierarchy Level [edit forwarding-options monitoring *group-name* family inet output]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the output interface for monitored traffic.

Options *interface-name*—Name of the interface.

 engine-id *number*—Identity of the monitoring interface.

 engine-type *number*—Type of this monitoring interface.

 input-interface-index *number*—Input interface index for records from this interface.

 output-interface-index *number*—Output interface index for records from this interface.

 source-address *address*—Address used for generating packets.

Usage Guidelines See “Configuring Flow Monitoring” on page 290.

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

interface (Next-Hop Group)

Syntax interface *interface-name* {
 next-hop [*addresses*];
 }

Hierarchy Level [edit forwarding-options next-hop-group *group-names*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the output interface for sending copies of packets elsewhere to be analyzed.

Options *interface-name*—Name of the interface.

 The remaining statements are explained separately.

Usage Guidelines See “Configuring a Next-Hop Group” on page 291.

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

interface (Port Mirroring)

Syntax	interface <i>interface-name</i> { next-hop <i>address</i> ; }
Hierarchy Level	[edit forwarding-options port-mirroring output], [edit forwarding-options port-mirroring family (inet inet6) output]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the output interface for sending copies of packets elsewhere to be analyzed.
Options	<i>interface-name</i> —Name of the interface. The remaining statements are explained separately.
Usage Guidelines	See “Configuring Port Mirroring” on page 330.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

layer2-unicast-replies

Syntax	layer2-unicast-replies;
Hierarchy Level	[edit forwarding-options dhcp-relay overrides], [edit forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Override the setting of the broadcast bit in DHCP request packets and instead use the Layer 2 unicast transmission method to transmit DHCP Offer reply packets and DHCP ACK reply packets from the DHCP server to DHCP clients during the discovery process.
Usage Guidelines	See “Configuring the Extended DHCP Relay Agent” on page 299.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

load-balance

Syntax	<pre>load-balance { <indexed-next-hop>; per-flow { hash-seed <i>number</i>; } per-prefix { hash-seed <i>number</i>; } }</pre>
Hierarchy Level	[edit forwarding-options], [edit logical-systems <i>logical-system-name</i> forwarding-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options], [edit routing-instances <i>routing-instance-name</i> forwarding-options]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Enable per-prefix or per-flow load balancing so that the router elects a next hop independently of the route selected by other routers.
Options	The statements are explained separately.
Usage Guidelines	See “Configuring Per-Prefix Load Balancing” on page 293 and “Configuring Per-Flow Load Balancing Based on Hash Values” on page 293.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	per-prefix

local-dump

Syntax	(local-dump no-local-dump);
Hierarchy Level	[edit forwarding-options sampling output cflowd <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable collection of cflowd records in a log file.
Options	<p>no-local-dump—Do not dump cflowd records to a log file before exporting.</p> <p>local-dump—Dump cflowd records to a log file before exporting.</p>
Usage Guidelines	See “Debugging cflowd Flow Aggregation” on page 328.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

local-server-group

Syntax	<code>local-server-group <i>local-server-group-name</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)]</p>
Release Information	Statement introduced in JUNOS Release 9.0.
Description	<p>Forward DHCP client packets to a specific extended DHCP local server when you use the DHCP vendor class identifier option (option 60) in DHCP packets to forward client traffic to specific DHCP servers.</p> <p>If the option 60 string received in the DHCP client packet matches the ASCII or hexadecimal match string and match criteria (exact match or partial match) that you specify, the extended DHCP relay agent forwards the client packets to the specified extended DHCP local server group configured with the <code>dhcp-local-server</code> statement at the [edit system services] hierarchy level.</p>
Options	<i>local-server-group-name</i> —Name of the extended DHCP local server group.
Usage Guidelines	See “Using Option 60 Information to Forward Client Traffic to Specific DHCP Servers” on page 306.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

logical-system-name

Syntax	logical-system-name;
Hierarchy Level	[edit forwarding-options dhcp-relay authentication username-include], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Specify that the logical system name be concatenated with the username during the subscriber authentication process. No logical system name is concatenated if the configuration is in the default logical system.
Usage Guidelines	See “Using External AAA Authentication Services” on page 313.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mac-address

Syntax	mac-address;
Hierarchy Level	[edit forwarding-options dhcp-relay authentication username-include], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Specify that the MAC address from the client PDU be concatenated with the username during the subscriber authentication process.
Usage Guidelines	See “Using External AAA Authentication Services” on page 313.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

maximum-capture-size

Syntax	maximum-capture-size <i>number</i> ;
Hierarchy Level	[edit forwarding-options packet-capture]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Configure the maximum size of capture for packets.
Options	<i>number</i> —Maximum capture size. Default: 68 bytes Range: 68 through 1500 bytes
Usage Guidelines	See “Configuring Packet Capture” on page 332.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

maximum-hop-count

Syntax	maximum-hop-count <i>number</i> ;
Hierarchy Level	[edit forwarding-options helpers bootp], [edit forwarding-options helpers bootp interface <i>interface-group</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	The maximum number of hops allowed.
Options	<i>number</i> —Maximum number of hops. Default: 4 hops
Usage Guidelines	See “Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent” on page 294.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

max-packets-per-second

Syntax	max-packets-per-second <i>number</i> ;
Hierarchy Level	[edit forwarding-options sampling input family]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify that the traffic threshold that must be exceeded before packets are dropped. A value of 0 instructs the Packet Forwarding Engine not to sample any traffic.
Options	<i>number</i> —Maximum number of packets per second. Range: 0 through 65,535 Default: 1000
Usage Guidelines	See “Configuring Traffic Sampling” on page 288.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

minimum-wait-time

Syntax	minimum-wait-time <i>seconds</i> ;
Hierarchy Level	[edit forwarding-options helpers bootp], [edit forwarding-options helpers bootp interface <i>interface-group</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	The minimum time allowed.
Options	<i>seconds</i> —Minimum time. Default: 3
Usage Guidelines	See “Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent” on page 294.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

monitoring

```

Syntax  monitoring group-name {
            family inet {
                output {
                    cflowd hostname {
                        port port-number;
                    }
                    export-format cflowd-version-5;
                    flow-active-timeout seconds;
                    flow-export-destination {
                        (cflowd-collector | collector-pic);
                    }
                    flow-inactive-timeout seconds;
                    interface interface-name {
                        engine-id number;
                        engine-type number;
                        input-interface-index number;
                        output-interface-index number;
                        source-address address;
                    }
                }
            }
        }

```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify flow monitoring instance name and properties.

The statements are explained separately.

Usage Guidelines See “Configuring Flow Monitoring” on page 290.

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

next-hop

Syntax	<code>next-hop <i>address</i>;</code>
Hierarchy Level	[edit forwarding-options port-mirroring output interface <i>interface-name</i>], [edit forwarding-options port-mirroring family (inet inet6) output interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the next-hop address for sending copies of packets to an analyzer.
Options	<i>address</i> —IP address of the next-hop router.
Usage Guidelines	See “Configuring Port Mirroring” on page 330.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

next-hop-group

Syntax	<code>next-hop-group [<i>group-names</i>] { interface <i>interface-name</i> { next-hop [<i>addresses</i>]; } }</code>
Hierarchy Level	[edit forwarding-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the next-hop address for sending copies of packets to an analyzer.
Options	<i>addresses</i> —IP address of the next-hop router. Each next-hop group supports up to 16 next-hop addresses. Up to 30 next-hop groups are supported. Each next-hop group must have at least two next-hop addresses. <i>group-names</i> —Name of next-hop group. Up to 30 next-hop groups are supported for the router. Each next-hop group must have at least 2 next-hop addresses. <i>interface-name</i> —Interface used to reach the next-hop destination.
Usage Guidelines	See “Configuring a Next-Hop Group” on page 291.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

no-filter-check

Syntax	no-filter-check;
Hierarchy Level	[edit forwarding-options port-mirroring output], [edit forwarding-options port-mirroring family (inet inet6) output]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable filter checking on the port-mirroring interface. This statement is required when you send port-mirrored traffic to a Tunnel PIC that has a filter applied to it.
Usage Guidelines	See “Configuring Port Mirroring” on page 330.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

no-listen

Syntax	no-listen;
Hierarchy Level	[edit forwarding-options helpers bootp interface <i>interface-group</i>], [edit forwarding-options helpers domain interface <i>interface-name</i>], [edit forwarding-options helpers tftp interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable recognition of DNS requests or stop packets from being forwarded on a logical interface, a group of logical interfaces, or a router.
Usage Guidelines	See “Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent” on page 294 and “Configuring DNS and TFTP Packet Forwarding” on page 295.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

no-local-dump

See local-dump

no-stamp

See stamp

no-world-readable

See world-readable**option-60**

Syntax option-60;

Hierarchy Level [edit forwarding-options dhcp-relay authentication username-include],
 [edit forwarding-options dhcp-relay group *group-name* authentication username-include],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay authentication username-include],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* authentication username-include],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay authentication username-include],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* authentication username-include],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay authentication username-include],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* authentication username-include]

Release Information Statement introduced in JUNOS Release 9.1.**Description** Specify that the payload of the Option 60 (Vendor Class Identifier) from the client PDU be concatenated with the username during the subscriber authentication process.**Usage Guidelines** See “Using External AAA Authentication Services” on page 313.

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

option-82

Syntax	option-82 [circuit-id] [remote-id];
Hierarchy Level	<p>[edit forwarding-options dhcp-relay authentication username-include],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]</p>
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Specify the type of option 82 information from the client PDU that is concatenated with the username during the subscriber authentication process. You can specify either, both, or neither the agent circuit ID nor the agent remote ID suboptions. If you specify both, the agent circuit ID is supplied first, followed by a delimiter, and then the agent remote ID. If you specify that neither suboption is supplied, the raw payload of option 82 from the PDU is concatenated to the username.
Options	<p>circuit-id—The string for the agent circuit ID suboption (suboption 1).</p> <p>remote-id—The string for the agent remote ID suboption (suboption 2).</p>
Usage Guidelines	See “Using External AAA Authentication Services” on page 313.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

output

See the following sections:

- output (Accounting) on page 401
- output (Forwarding Table) on page 402
- output (Monitoring) on page 402
- output (Port Mirroring) on page 403
- output (Sampling) on page 404

output (Accounting)

Syntax

```
output {
  cflowd [ hostnames ] {
    aggregation {
      autonomous-system;
      destination-prefix;
      protocol-port;
      source-destination-prefix {
        caida-compliant;
      }
      source-prefix;
    }
    autonomous-system-type (origin | peer);
    port port-number;
    version format;
  }
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
  interface[ interface-names ] {
    engine-id number;
    engine-type number;
    source-address address;
  }
}
```

Hierarchy Level [edit forwarding-options accounting *group-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure cflowd, output interfaces, and flow properties.

The statements are explained separately.

Usage Guidelines See “Configuring Discard Accounting” on page 289.

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

output (Forwarding Table)

Syntax	<code>output <i>filter-name</i>;</code>
Hierarchy Level	[edit forwarding-options family (inet inet6 mpls) filter], [edit routing-instance <i>instance-name</i> forwarding-options family (inet inet6 mpls) filter]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Configure filtering on the egress traffic of the forwarding table.
Options	<i>filter-name</i> —Name of the applied filter.
Usage Guidelines	See “Configuring a Forwarding Table Filter” on page 286.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

output (Monitoring)

Syntax	<pre>output { cflowd <i>hostname</i> { port <i>port-number</i>; } export-format cflowd-version-5; flow-active-timeout <i>seconds</i>; flow-export-destination { (cflowd-collector collector-pic); } flow-inactive-timeout <i>seconds</i>; interface <i>interface-name</i> { engine-id <i>number</i>; engine-type <i>number</i>; input-interface-index <i>number</i>; output-interface-index <i>number</i>; source-address <i>address</i>; } }</pre>
Hierarchy Level	[edit forwarding-options monitoring <i>group-name</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure cflowd, output interfaces, and flow properties. The statements are explained separately.
Usage Guidelines	See “Configuring Flow Monitoring” on page 290.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

output (Port Mirroring)

Syntax output {
 interface *interface-name* {
 next-hop *address*;
 }
 no-filter-check;
 }

Hierarchy Level [edit forwarding-options port-mirroring family (inet | inet6)]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure output interfaces and flow properties.

The statements are explained separately.

Usage Guidelines See “Configuring Port Mirroring” on page 330.

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

output (Sampling)

Syntax

```

output {
  cflowd hostname {
    version9 {
      template template-name;
    }
    aggregation {
      autonomous-system;
      destination-prefix;
      protocol-port;
      source-destination-prefix {
        caida-compliant;
      }
      source-prefix;
    }
    autonomous-system-type (origin | peer);
    (local-dump | no-local-dump);
    port port-number;
    source-address address;
    version format;
  }
  file {
    disable;
    filename filename;
    files number;
    size (Sampling) bytes;
    (stamp | no-stamp);
    (world-readable | no-world-readable);
  }
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
  interface [ interface-names ] {
    engine-id number;
    engine-type number;
    source-address address;
  }
}

```

Hierarchy Level [edit forwarding-options sampling]

Release Information Statement introduced before JUNOS Release 7.4.
version9 statement introduced in JUNOS Release 8.3.

Description Configure cflowd, output files and interfaces, and flow properties. Enable the collection of traffic flows using the version 9 format.

The statements are explained separately.

Usage Guidelines See “Configuring Traffic Sampling Output” on page 324 and “Configuring Active Flow Monitoring Using Version 9” on page 329.

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

overrides

Syntax	<pre> overrides { always-write-giaddr; always-write-option-82; layer2-unicast-replies; trust-option-82; disable-relay; } </pre>
Hierarchy Level	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>] </pre>
Release Information	Statement introduced in JUNOS Release 8.3.
Description	<p>Override the default configuration settings for the extended DHCP relay agent. Specifying the overrides statement with no subordinate statements removes all DHCP relay agent overrides at that hierarchy level.</p> <p>To override global DHCP relay agent configuration options, include the overrides statement and its subordinate statements at the [edit forwarding-options dhcp-relay] hierarchy level. To override DHCP relay agent configuration options for a named group of interfaces, include the statements at the [edit forwarding-options dhcp-relay group <i>group-name</i>] hierarchy level.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring the Extended DHCP Relay Agent” on page 299.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

packet-capture

Syntax packet-capture {
 disable;
 file {
 filename *file-name*;
 files *number*;
 size *number*;
 }
 maximum-capture-size *number*;
 }

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced in JUNOS Release 7.5.

Description Configure packet capture on a router.

Options disable—Disable packet capture on the router.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Packet Capture” on page 332.

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

password

Syntax	<code>password <i>password-string</i>;</code>
Hierarchy Level	[edit forwarding-options dhcp-relay authentication], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Configure the password that is sent to the external AAA authentication server for subscriber authentication.
Options	<i>password-string</i> —Authentication password.
Usage Guidelines	See “Using External AAA Authentication Services” on page 313.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

per-flow

Syntax	per-flow { hash-seed <i>number</i> ; }
Hierarchy Level	[edit forwarding-options load-balance], [edit logical-systems <i>logical-system-name</i> forwarding-options load-balance], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options load-balance], [edit routing-instances <i>routing-instance-name</i> forwarding-options load-balance]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Configure the hash parameter for per-flow load balancing.
Options	hash-seed—Configure the hash value. <i>number</i> —Hash value. Range: 0 through 65,535 Default: 0
Usage Guidelines	See “Configuring Per-Flow Load Balancing Based on Hash Values” on page 293.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	load-balance

per-prefix

Syntax	per-prefix { hash-seed <i>number</i> ; }
Hierarchy Level	[edit forwarding-options load-balance], [edit logical-systems <i>logical-system-name</i> forwarding-options load-balance], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options load-balance], [edit routing-instances <i>routing-instance-name</i> forwarding-options load-balance]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Configure the hash parameter for per-prefix load balancing.
Options	hash-seed—Configure the hash value. <i>number</i> —Hash value. Range: 0 through 65,535 Default: 0
Usage Guidelines	See “Configuring Per-Prefix Load Balancing” on page 293.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	load-balance

port

Syntax	port <i>port-number</i> ;
Hierarchy Level	[edit forwarding-options accounting <i>group-name</i> output cflowd <i>hostname</i>], [edit forwarding-options monitoring <i>group-name</i> family inet output cflowd <i>hostname</i>], [edit forwarding-options sampling output cflowd <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the UDP port number on the cflowd host system.
Options	<i>port-number</i> —Any valid UDP port number on the host system.
Usage Guidelines	See “Configuring Flow Aggregation (cflowd)” on page 326.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

port-mirroring

Syntax

```
port-mirroring {
  family (inet | inet6) {
    input {
      rate number;
      run-length number;
    }
    output {
      interface interface-name {
        next-hop address;
      }
      no-filter-check;
    }
  }
  traceoptions {
    file filename {
      files number;
      size bytes;
      (world-readable | no-world-readable);
    }
  }
}
```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the address family, rate, run length, interface, and next-hop address for sending copies of packets to an analyzer.

The statements are explained separately.



NOTE: For information about configuring port mirroring for Layer 2 traffic on MX-series routers, see the *JUNOS MX-series Layer 2 Configuration Guide*.

Usage Guidelines See “Configuring Port Mirroring” on page 330.

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

prefix

Syntax	prefix host-name logical-system-name routing-instance-name;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay relay-option-82 circuit-id],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> relay-option-82 circuit-id],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay relay-option-82 circuit-id],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-82 circuit-id],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-82 circuit-id],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-82 circuit-id],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-82 circuit-id],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-82 circuit-id]</p>
Release Information	Statement introduced in JUNOS Release 8.3.
Description	<p>Add a prefix to the base option 82 agent-circuit-id information in DHCP packets destined for a DHCP server. The prefix can consist of any combination of the hostname, logical system name, and routing instance name.</p> <p>If you include only the hostname, only the logical system name, or only the routing instance name in the prefix, the format of the agent-circuit-id information for Fast Ethernet or Gigabit Ethernet interfaces with stacked virtual LANs (S-VLANs) is one of the following:</p> <pre> host-name:(fe ge)-fpc/pic/port:svlan-id-vlan-id logical-system-name:(fe ge)-fpc/pic/port:svlan-id-vlan-id routing-instance-name:(fe ge)-fpc/pic/port:svlan-id-vlan-id </pre> <p>If you include both the logical system name and the routing instance name in the prefix, the format of the agent-circuit-id information for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs is as follows:</p> <pre> logical-system-name;routing-instance-name:(fe ge)-fpc/pic/port:svlan-id-vlan-id </pre> <p>If you include the hostname, logical system name, and routing instance name in the prefix, the format of the agent-circuit-id information for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs is as follows:</p> <pre> host-name/logical-system-name;routing-instance-name:(fe ge)-fpc/pic/port:svlan-id-vlan-id </pre> <p>For Fast Ethernet or Gigabit Ethernet interfaces that use virtual LANs (VLANs) but not S-VLANs, only the <i>vlan-id</i> value appears in the agent-circuit-id format. For Fast Ethernet or Gigabit Ethernet interfaces that do not use VLANs or S-VLANs, neither the <i>vlan-id</i> value nor the <i>svlan-id</i> value appears.</p>

Options **host-name**—Prepend the hostname of the router configured with the **host-name** statement at the [edit system] hierarchy level to the agent-circuit-id information.

logical-system-name—Prepend the name of the logical system to the agent-circuit-id information.

routing-instance-name—Prepend the name of the routing instance to the agent-circuit-id information.

Usage Guidelines See “Configuring the Extended DHCP Relay Agent” on page 299.

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

rate

Syntax *rate number;*

Hierarchy Level [edit forwarding-options port-mirroring family (inet | inet6) input],
[edit forwarding-options port-mirroring input family inet],
[edit forwarding-options sampling input family]

Release Information Statement introduced before JUNOS Release 7.4.

Description Set the ratio of the number of packets to be sampled. For example, if you specify a rate of 10, every tenth packet (1 packet out of 10) is sampled.

Options *number*—Denominator of the ratio.
Range: 1 through 65,535

Usage Guidelines See “Configuring Traffic Sampling” on page 288.

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

relay-option-60

Syntax	<pre> relay-option-60 { vendor-option { (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>) { (relay-server-group <i>server-group-name</i> local-server-group <i>local-server-group-name</i> drop); } (default-relay-server-group <i>server-group-name</i> default-local-server-group <i>local-server-group-name</i> drop); } }</pre>
Hierarchy Level	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>]</pre>
Release Information	Statement introduced in JUNOS Release 9.0.
Description	<p>Configure the extended DHCP relay agent to use the DHCP vendor class identifier option (option 60) in DHCP client packets to forward client traffic to specific DHCP servers, or to drop selected DHCP client packets. This feature is useful in network environments where DHCP clients access services provided by multiple vendors and DHCP servers.</p> <p>You can use the relay-option-60 statement and its subordinate statements at the [edit forwarding-options dhcp-relay] hierarchy level to configure option 60 support globally, or at the [edit forwarding-options dhcp-relay group <i>group-name</i>] hierarchy level to configure option 60 support for a named group of interfaces. You can also configure option 60 support for the extended DHCP relay agent on a per logical system and per routing instance basis.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Using Option 60 Information to Forward Client Traffic to Specific DHCP Servers” on page 306.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

relay-option-82

Syntax relay-option-82 {
 circuit-id {
 prefix host-name logical-system-name routing-instance-name;
 }
 }

Hierarchy Level [edit forwarding-options dhcp-relay],
 [edit forwarding-options dhcp-relay group *group-name*],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay group *group-name*],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group
 group-name]

Release Information Statement introduced in JUNOS Release 8.3.

Description Enable or disable the insertion of the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server.

If you enable insertion of option 82 information in DHCP packets, you must specify at least the **circuit-id** statement to include the agent-circuit-id suboption (suboption 1) of the DHCP relay agent information option. Optionally, you can also specify the **prefix** statement to add a prefix to the base option 82 information that consists of any combination of the hostname, logical system name, and routing instance name. Specifying the **relay-option-82** statement with no subordinate statements disables insertion of option 82 information in DHCP packets, which is the default behavior.

You can use the **relay-option-82** statement and its subordinate statements at the [edit forwarding-options dhcp-relay] hierarchy level to control insertion of option 82 information globally, or at the [edit forwarding-options dhcp-relay group *group-name*] hierarchy level to control insertion of option 82 information for a named group of interfaces.

The statements are explained separately.

Usage Guidelines See “Configuring the Extended DHCP Relay Agent” on page 299.

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

relay-server-group

Syntax	<code>relay-server-group server-group-name;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-60 vendor-option (equals starts-with) (ascii <i>match-string</i> hexadecimal <i>match-hex</i>)]</p>
Release Information	Statement introduced in JUNOS Release 9.0.
Description	<p>Relay DHCP client packets to the specified group of extended DHCP relay servers when you use the DHCP vendor class identifier option (option 60) in DHCP packets to forward client traffic to specific DHCP servers.</p> <p>If the option 60 string received in the DHCP client packet matches the ASCII or hexadecimal match string and match criteria (exact match or partial match) that you specify, the extended DHCP relay agent relays the client packets to the specified group of servers configured with the server-group statement at the [edit forwarding-options dhcp-relay] hierarchy level. A server group can contain multiple server addresses and can map to more than one ASCII or hexadecimal match string.</p>
Options	<i>server-group-name</i> —Name of the extended DHCP relay server group.
Usage Guidelines	See “Using Option 60 Information to Forward Client Traffic to Specific DHCP Servers” on page 306.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

route-accounting

Syntax	route-accounting;
Hierarchy Level	[edit forwarding-options family inet6]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Configure the routing platform to track IPv6 traffic passing through the router.
Default	Disabled
Usage Guidelines	See “Configuring IPv6 Accounting” on page 287.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

routing-instance

Syntax	routing-instance [<i>routing-instance-names</i>];
Hierarchy Level	[edit forwarding-options helpers domain interface <i>interface-name</i> server <i>address</i>], [edit forwarding-options helpers domain server <i>address</i>], [edit forwarding-options helpers tftp interface <i>interface-name</i> server <i>address</i>], [edit forwarding-options helpers tftp server <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the routing instance of the server to forward. You can include as many routing instances as necessary in the same statement.
Options	<i>routing-instance-names</i> —Name of the routing instance.
Usage Guidelines	See “Configuring a Forwarding Table Filter” on page 286.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

routing-instance-name

Syntax	routing-instance-name;
Hierarchy Level	[edit forwarding-options dhcp-relay authentication username-include], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Specify that the routing instance name be concatenated with the username during the subscriber authentication process. No routing instance name be concatenated if the configuration is in the default routing instance.
Usage Guidelines	See “Using External AAA Authentication Services” on page 313.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

run-length

Syntax	run-length <i>number</i> ;
Hierarchy Level	[edit forwarding-options port-mirroring family (inet inet6) input], [edit forwarding-options sampling input family]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the number of samples following the initial trigger event. This allows you to sample packets following those already being sampled.
Options	<i>number</i> —Number of samples. Range: 0 through 20 Default: 0
Usage Guidelines	See “Configuring a Forwarding Table Filter” on page 286.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

sampling

```

Syntax  sampling {
    disable;
    input {
        family (inet | mpls) {
            max-packets-per-second number;
            rate number;
            run-length number;
        }
    }
    output {
        cflowd hostname {
            version9 {
                template template-name;
            }
        }
        aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
                caida-compliant;
            }
            source-prefix;
        }
        autonomous-system-type (origin | peer);
        (local-dump | no-local-dump);
        port port-number;
        source-address address;
        version format;
    }
    file {
        disable;
        filename filename;
        files number;
        size bytes;
        (stamp | no-stamp);
        (world-readable | no-world-readable);
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface [ interface-names ] {
        engine-id number;
        engine-type number;
        source-address address;
    }
}
traceoptions {
    file filename {
        files number;
        size bytes;
        (world-readable | no-world-readable);
    }
}

```



```
}
}
```

Hierarchy Level	[edit forwarding-options]
Release Information	Statement introduced before JUNOS Release 7.4. sampling statement introduced in JUNOS Release 8.3.
Description	Configure traffic sampling. The statements are explained separately.
Usage Guidelines	See “Configuring a Forwarding Table Filter” on page 286, “Configuring Flow Aggregation (cflowd)” on page 326, “Tracing Traffic Sampling Operations” on page 326, “Configuring Active Flow Monitoring Using Version 9” on page 329, and “Configuring Port Mirroring” on page 330.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

server

See the following sections:

- server (DHCP or BOOTP Service) on page 420
- server (DNS and TFTP Service) on page 420

server (DHCP or BOOTP Service)

Syntax	server [<i>addresses</i>];
Hierarchy Level	[edit forwarding-options helpers bootp], [edit forwarding-options helpers bootp interface <i>interface-group</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the router to act as a DHCP or BOOTP relay agent.
Options	<i>addresses</i> —One or more addresses of the server.
Usage Guidelines	See “Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent” on page 294.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

server (DNS and TFTP Service)

Syntax	server <i>address</i> < [routing-instance <i>routing-instance-names</i>] >;
Hierarchy Level	[edit forwarding-options helpers domain], [edit forwarding-options helpers domain interface <i>interface-name</i>], [edit forwarding-options helpers tftp], [edit forwarding-options helpers tftp interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the DNS or TFTP server for forwarding DNS or TFTP requests. Only one server can be specified for each interface.
Options	<i>address</i> —Address of the server. routing-instance [<i>routing-instance-names</i>]—Set the routing instance name or names that belong to the DNS server or TFTP server.
Usage Guidelines	See “Configuring DNS and TFTP Packet Forwarding” on page 295.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

server-group

Syntax	<pre>server-group { server-group-name { server-ip-address; } }</pre>
Hierarchy Level	[edit forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Specify the name of a group of DHCP server addresses for use by the extended DHCP relay agent.
Options	<p><i>server-group-name</i>—Name of the group of DHCP server addresses.</p> <p><i>server-ip-address</i>—IP address of the DHCP server belonging to this named server group. You can configure a maximum of five IP addresses per named server group.</p>
Usage Guidelines	See “Configuring the Extended DHCP Relay Agent” on page 299.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

size

See the following sections:

- [size \(Packet Capture\)](#) on page 422
- [size \(Sampling\)](#) on page 423

size (Packet Capture)

Syntax *size number;*

Hierarchy Level [edit forwarding-options packet-capture file]

Release Information Statement introduced in JUNOS Release 7.5.

Description Configure the maximum size of the file for packet capturing.

Options *number*—Maximum size of file.

Range: 1024 through 104,857,600 bytes

Default: 512,000 bytes

Usage Guidelines See “Configuring Packet Capture” on page 332.

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

size (Sampling)

Syntax	size bytes;
Hierarchy Level	[edit forwarding-options port-mirroring traceoptions file], [edit forwarding-options sampling output file], [edit forwarding-options sampling traceoptions file]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the maximum size of each file containing sample or log data. The file size is limited by the number of files to be created and the available hard disk space. When a traffic sampling file named sampling-file reaches the maximum size, it is renamed sampling-file.0 . When the sampling-file file again reaches its maximum size, sampling-file.0 is renamed sampling-file.1 and sampling-file is renamed sampling-file.0 . This renaming scheme continues until the maximum number of traffic sampling files is reached. Then the oldest traffic sampling file is overwritten.
Options	bytes —Maximum size of each traffic sampling file or trace log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). Syntax: xk to specify KB, xm to specify MB, or xg to specify GB Range: 10 KB through the maximum file size supported on your router Default: 1 MB for sampling data; 128 KB for log information
Usage Guidelines	See “Configuring Traffic Sampling Output” on page 324 and “Tracing Traffic Sampling Operations” on page 326.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

stamp

Syntax	(stamp no-stamp);
Hierarchy Level	[edit forwarding-options sampling output file]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Include a timestamp with each line in the output file.
Options	no-stamp —Do not include timestamps. This is the default. stamp —Include a timestamp with each line of packet sampling information. Default: No timestamp is included.
Usage Guidelines	See “Configuring Traffic Sampling Output” on page 324.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

tftp

Syntax tftp {
 description *description-of-service*;
 server address < [routing-instance *routing-instance-names*] >;
 interface *interface-name* {
 description *description-of-interface*;
 no-listen;
 server address < [routing-instance *routing-instance-names*] >;
 }
 }

Hierarchy Level [edit forwarding-options helpers]

Release Information Statement introduced before JUNOS Release 7.4.

Description Enable TFTP request packet forwarding.

Options The statements are explained separately.

Usage Guidelines See “Configuring DNS and TFTP Packet Forwarding” on page 295.

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

traceoptions

See the following sections:

- `traceoptions` (DNS and TFTP Packet Forwarding) on page 426
- `traceoptions` (Extended DHCP Relay Agent) on page 428
- `traceoptions` (Port Mirroring and Traffic Sampling) on page 430

traceoptions (DNS and TFTP Packet Forwarding)

Syntax traceoptions {
 file *filename* <files *number*> <size *size*> <world-readable | no-world-readable> <match
 regex>;
 flag *flag*;
 }

Hierarchy Level [edit forwarding-options helpers]

Release Information Statement introduced before JUNOS Release 7.4.
 Statement standardized in JUNOS Release 8.0.
 match option introduced in JUNOS Release 8.0.

Description Configure tracing operations.

Default If you do not include this statement, no tracing operations are performed.

Options *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in a file named **fud** in the directory **/var/log**. If you include the **file** statement, you must specify a filename.

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 also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000

Default: 3 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Log all operations
- **configuration**—Log reading of configuration at the [edit event-options] hierarchy level
- **events**—Log fud processing
- **database**—Log events involving storage and retrieval in events database
- **policy**—Log policy processing
- **server**—Log communication with processes that are generating events
- **timer-events**—Log internally generated events

match regex—(Optional) Refine the output to include lines that contain the regular expression.

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 the **trace-file** file again reaches its maximum 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 also must specify a maximum number of trace files with the `files` option and filename.

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

Range: 0 bytes through 4,294,967,295 KB

Default: 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

The remaining statements are explained separately.

Usage Guidelines See “Tracing BOOTP, DNS, and TFTP Forwarding Operations” on page 296.

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

traceoptions (Extended DHCP Relay Agent)

Syntax	<pre> traceoptions { file <i>file-name</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable> <match <i>regex</i>>; flag <i>flag</i>; } </pre>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay]</p>
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Configure tracing operations for extended DHCP relay agent processes.
Default	If you do not include this statement, no tracing operations are performed.
Options	<p><i>file-name</i>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (“ ”). All files are placed in a file named <i>jdhcpd</i> in the directory <i>/var/log</i>. If you include the file statement, you must specify a filename.</p> <p><i>files number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option and a filename.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p><i>flag flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> ■ all—Trace all events ■ auth—Trace authentication events ■ database—Trace database events ■ fwd—Trace firewall process events ■ general—Trace miscellaneous events ■ ha—Trace high availability-related events ■ interface—Trace interface operations ■ io—Trace I/O operations ■ packet—Trace packet decoding operations ■ packet-option—Trace DHCP option decoding operations ■ rpd—Trace routing protocol process events

- **rtsock**—Trace routing socket operations
- **session-db**—Trace session database operations
- **state**—Trace changes in state
- **ui**—Trace user interface operations

match *regex*—(Optional) Refine the output to include lines that contain the regular expression.

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 the **trace-file** file again reaches its maximum 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 also must specify a maximum number of trace files with the **files** option and filename.

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

Range: 0 bytes through 4,294,967,295 KB

Default: 128 KB

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

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Extended DHCP Relay Agent” on page 299.

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

traceoptions (Port Mirroring and Traffic Sampling)

Syntax	<pre> traceoptions { file <i>filename</i> { files <i>number</i>; size <i>bytes</i>; (world-readable no-world-readable); } } </pre>
Hierarchy Level	[edit forwarding-options port-mirroring], [edit forwarding-options sampling]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Configure traffic sampling tracing operations.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Tracing Traffic Sampling Operations” on page 326.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

trust-option-82

Syntax	trust-option-82;
Hierarchy Level	<pre> [edit forwarding-options dhcp-relay overrides], [edit forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides] </pre>
Release Information	Statement introduced in JUNOS Release 8.3.
Description	<p>Enable processing of DHCP client packets that have a gateway IP address (giaddr) of 0 (zero) and contain option 82 information. By default, the DHCP relay agent treats such packets as if they originated at an untrusted source, and drops them without further processing.</p>
Usage Guidelines	See “Configuring the Extended DHCP Relay Agent” on page 299.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

username-include

Syntax username-include {
 circuit-type;
 delimiter *delimiter-character*;
 domain-name *domain-name-string*;
 logical-system-name;
 mac-address;
 option-60;
 option-82 [circuit-id] [remote-id];
 routing-instance-name;
 user-prefix *user-prefix-string*;
 }

Hierarchy Level [edit forwarding-options dhcp-relay authentication],
 [edit forwarding-options dhcp-relay group *group-name* authentication],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay authentication],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name*
 authentication],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay authentication],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay group *group-name* authentication],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay authentication],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group
 group-name authentication]

Release Information Statement introduced in JUNOS Release 9.1.

Description Configure the username that the router passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router accesses the local authentication service only and does not use external authentication services, such as RADIUS.

The statements are explained separately.

Usage Guidelines See “Using External AAA Authentication Services” on page 313.

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

user-prefix

Syntax	<code>user-prefix <i>user-prefix-string</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay authentication username-include],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]</p>
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Specify the user prefix that is concatenated with the username during the subscriber authentication process.
Options	<i>user-prefix-string</i> —The user prefix string.
Usage Guidelines	See “Using External AAA Authentication Services” on page 313.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

vendor-option

Syntax vendor-option {
 (equals | starts-with) (ascii *match-string* | hexadecimal *match-hex*) {
 (relay-server-group *server-group-name* | local-server-group *local-server-group-name* |
 drop);
 }
 (default-relay-server-group *server-group-name* | default-local-server-group
 local-server-group-name | drop);
 }

Hierarchy Level [edit forwarding-options dhcp-relay relay-option-60],
 [edit forwarding-options dhcp-relay group *group-name* relay-option-60],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay relay-option-60],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name*
 relay-option-60],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay relay-option-60],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay group *group-name* relay-option-60],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay relay-option-60],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group
group-name relay-option-60]

Release Information Statement introduced in JUNOS Release 9.0.

Description Configure the match criteria when you use the DHCP vendor class identifier option (option 60) in DHCP client packets to forward client traffic to specific DHCP servers. The extended DHCP relay agent compares the option 60 vendor-specific strings received in DHCP client packets against the match criteria that you specify. If there is a match, you can define certain actions for the associated DHCP client packets.

The **vendor-option** statement enables you to specify either an exact, left-to-right match (with the **equals** statement) or a partial match (with the **starts-with** statement), and configure either an ASCII match string (with the **ascii** statement) or a hexadecimal match string (with the **hexadecimal** statement).

You can configure an unlimited number of match strings. Match strings do not support the use of wildcard attributes.

Options **equals**—Exact, left-to-right match of the ASCII or hexadecimal match string with the option 60 string.

starts-with—Partial match of the ASCII or hexadecimal match string with the option 60 string. The option 60 string can contain a superset of the ASCII or hexadecimal match string, provided that the leftmost characters of the option 60 string entirely match the characters in the configured match string. When you use the **starts-with** statement, the longest match rule applies; that is, the router matches the string “test123” before it matches the string “test”.

ascii *match-string*—ASCII match string of 1 through 255 alphanumeric characters.

hexadecimal *match-hex*—Hexadecimal match string of 1 through 255 hexadecimal characters (0 through 9, a through f, A through F).

The remaining statements are explained separately.

Usage Guidelines See “Using Option 60 Information to Forward Client Traffic to Specific DHCP Servers” on page 306.

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

version

Syntax *version format;*

Hierarchy Level [edit forwarding-options accounting *group-name* output cflowd *hostname*],
[edit forwarding-options sampling output cflowd *hostname*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the version format of the aggregated flows exported to a cflowd server.

Options *format*—Export format of the flows.

Values: 5 or 8

Default: 5

Usage Guidelines See “Configuring Flow Aggregation (cflowd)” on page 326.

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

version9

Syntax	version9 { template <i>template-name</i> ; }
Hierarchy Level	[edit forwarding-options sampling cflowd <i>hostname</i>]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Enable active flow monitoring using the version 9 template format to collect traffic flows.
Options	<i>template-name</i> —Name of a version 9 record flow format template configured at the [edit services monitoring] hierarchy level.
Usage Guidelines	See “Configuring Active Flow Monitoring Using Version 9” on page 329
Required Privilege Level	interface—To view this statement in the configuration. interface-level—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i> and <i>JUNOS Feature Guide</i>

world-readable

Syntax	(world-readable no-world-readable);
Hierarchy Level	[edit forwarding-options port-mirroring traceoptions file], [edit forwarding-options sampling output file], [edit forwarding-options sampling traceoptions file]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable unrestricted file access.
Options	no-world-readable—Restrict file access to owner. This is the default. world-readable—Enable unrestricted file access. Default: no-world-readable
Usage Guidelines	See “Configuring Traffic Sampling Output” on page 324 and “Tracing Traffic Sampling Operations” on page 326.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

Part 5

Indexes

- Index on page 439
- Index of Statements and Commands on page 449

Index

Symbols

!	
in policy expressions	
logical operator.....	60
! (negation)	
in firewall filters	
bit-field logical operator.....	190
#, comments in configuration statements.....	xxix
&&, logical operator.....	60
&, bit-field logical operator.....	190
(), in syntax descriptions.....	xxix
+	
bit-field logical operator.....	190
, (comma), bit-field logical operator.....	190
< >, in syntax descriptions.....	xxix
[], in configuration statements.....	xxix
{ }, in configuration statements.....	xxix
(pipe)	
in firewall filters	
bit-field logical operator.....	190
(pipe), in syntax descriptions.....	xxix
(pipes), logical operator.....	60

A

accept	
firewall filters	
action.....	168, 171, 285
policy, routing	
control action.....	48
access and access-internal routes	
importing and exporting in routing policies.....	57
accounting statement.....	336
usage guidelines.....	289
accounting-profile statement	
firewall.....	257
usage guidelines.....	221
action modifiers, firewall filters.....	168, 171
action statement	
policer (TCM).....	258
usage guidelines.....	251

actions

firewall filters.....	160, 168
policy, routing.....	25
characteristics, manipulating.....	49
flow control.....	47, 48
tracing.....	47
tracing.....	53
active-server-group statement.....	337
usage guidelines.....	311
address (filter match conditions).....	184
address (firewall filter match condition).....	186
aggregate policer.....	246
aggregation statement.....	338
usage guidelines.....	289
always-write-giaddr statement.....	339
usage guidelines.....	304
always-write-option-82 statement.....	340
usage guidelines.....	304
ampersand (&), bit-field logical operator.....	190
apply-path statement.....	143
usage guidelines.....	111, 112
area (routing policy match condition).....	42
as-path (routing policy match condition).....	42
as-path statement.....	144
policy, routing	
usage guidelines.....	92
as-path-group statement.....	144
usage guidelines.....	92
as-path-prepend (routing policy action).....	49, 129
ASs	
paths	
modifying with routing policy.....	49, 129
regular expressions <i>See</i> policy, routing, AS path	
regular expressions	
authentication statement.....	341
autonomous-system-type statement.....	342
usage guidelines.....	327

B

bandwidth policer.....	247
------------------------	-----

BGP	
communities	
description	97
names	100
policy, routing	98, 145
damping parameters	130, 148
policy, routing	
applying	27
bgp extended community statement	
usage guidelines	107
bit-field	
firewall filter match conditions	188, 189
logical operators	190
BOOTP relay agent	294
bootp statement	343
usage guidelines	294
braces, in configuration statements	xxix
brackets	
angle, in syntax descriptions	xxix
square, in configuration statements	xxix
C	
cflowd statement	344
usage guidelines	326
chassis	
show chassis hardware command	277
circuit-id statement	347
usage guidelines	308
circuit-type statement	348
class (routing policy action)	49, 130
class-based firewall filter match conditions	190
client-response-ttl statement	348
color	
policy, routing	
action	49
match condition	42
comments, in configuration statements	xxix
communities	
description	97
extend range of BGP communities	107
names	100
policy, routing	98, 145
action	49
match condition	42
community statement	145
policy, routing	
usage guidelines	98
usage guidelines	107
condition statement	147
conventions	
text and syntax	xxviii
CoS	
modifying with routing policy	130
count (firewall filter action)	168, 171
curly braces, in configuration statements	xxix
customer support	xxxvii
contacting JTAC	xxxvii
D	
damping	
policy, routing, action	50
damping statement	148
BGP	
usage guidelines	134
policy, routing	
usage guidelines	131
default-local-server-group statement	349
usage guidelines	306
default-relay-server-group statement	350
usage guidelines	306
delimiter statement	351
description statement	
interface	352
usage guidelines	294, 295
service	352
usage guidelines	294, 295
destination class usage	50, 51
destination-address (firewall filter match	
condition)	186
destination-class (routing policy action)	50, 53
destination-port (firewall filter match	
condition)	176, 191
DHCP	
relay agents, DHCP/BOOTP	294
relay agents, extended	299
access and access-internal routes	57, 303
active server groups	311
clearing client address bindings and	
statistics	314
configuration examples	317, 318, 319, 320
graceful Routing Engine switchover	304
how components interact	303, 304
interface groups	312
option 82 information	308
option 60 information	306
server groups	311
state persistence	304
tracing operations	314
verifying client address bindings and	
statistics	314
dhcp-relay statement	353
usage guidelines	299
disable statement	356
traffic sampling	
usage guidelines	321
disable-relay statement	357
usage guidelines	304
discard (firewall filter action)	171
discard interface	87

DNS	
packet forwarding.....	295
requests, disabling recognition.....	295
documentation set	
comments on.....	xxxvii
domain statement.....	357
domain-name statement.....	358
drop statement.....	359
usage guidelines.....	306
dscp (firewall filter match condition).....	179
DVMRP	
policy, routing	
applying.....	27

E

ES-IS.....	332
evaluation	
firewall filters.....	173
policy, routing	28
event policy	
all (tracing flag).....	298
configuration (tracing flag).....	298
database (tracing flag).....	298
events (tracing flag).....	298
policy (tracing flag).....	298
server (tracing flag).....	298
timer-events (tracing flag).....	298
exact route list match type.....	113, 115
examples	
extended DHCP relay agent.....	317, 318, 319, 320
except (firewall filter match condition).....	176, 186
exclamation point (!), bit-field logical operator.....	190
export routing policies	
applying.....	26, 58, 59, 149
from statement.....	66
export statement.....	149
policy, routing	
usage guidelines.....	26, 58, 64
export-format statement.....	360
usage guidelines.....	290

F

family inet statement (firewall filter)	
usage guidelines.....	163
family inet statement (load balancing).....	363
usage guidelines.....	281
family mpls statement.....	365
usage guidelines.....	138
family multiservice statement.....	366
usage guidelines.....	140
family statement.....	361
firewall filter.....	259
usage guidelines.....	166

file statement	
helpers trace options.....	367
packet capture.....	368
sampling.....	368
traceoptions.....	369
traffic sampling output	
usage guidelines.....	296, 324, 326
filename statement.....	370
files	
firewall log output file.....	172
logging information output file.....	296, 326
traffic sampling output files.....	324
var/log/sampled file.....	296, 326
var/tmp/sampled.pkts file.....	324
files statement	
packet capture.....	371
sampling.....	371
usage guidelines.....	324
filter statement	
firewall.....	260
usage guidelines.....	166
forwarding table.....	372
usage guidelines.....	164
filter-specific statement.....	261
usage guidelines.....	239
filters	
interface-specific counters.....	206
firewall filters	
actions.....	160, 168
applying.....	205, 225
architecture	6, 8
comparison with routing policies	8, 11
configuration statements.....	163, 257
evaluation.....	173
example filter definitions.....	194, 207, 252
family address type.....	166
flow, packets	3
in traffic sampling.....	285
log output file.....	172
match conditions.....	160, 167, 174
names.....	166
numeric match conditions, IPv6 (table).....	176, 179
numeric match conditions, Layer 2 bridging	
(MX-series routers).....	182
overview.....	159, 231, 277
policing.....	233
purpose.....	8
service filters.....	203
show firewall filters command.....	171
simple filters.....	203
system logging.....	227
terms.....	166
testing packet protocols.....	191
firewall log output file.....	172
firewall statement.....	261
usage guidelines.....	163

first-fragment (firewall filter match condition).....	190
flood statement.....	372
usage guidelines.....	287
flooding	
IS-IS and OSPF.....	21
flow aggregation.....	326
flow control actions.....	47, 48
flow-active-timeout statement.....	373
accounting	
usage guidelines.....	290
sampling	
usage guidelines.....	324
flow-export-destination statement.....	373
usage guidelines.....	291
flow-inactive-timeout statement.....	374
accounting	
usage guidelines.....	290
sampling	
usage guidelines.....	324
font conventions.....	xxviii
forwarding table	
filters.....	225, 286
policy, routing	
applying.....	28, 66
forwarding-class (firewall filter action).....	168, 171
forwarding-options statement.....	374
usage guidelines.....	279
fragment-flags (firewall filter match condition).....	190
fragment-offset (firewall filter match condition).....	176
from statement.....	153
firewall filters	
usage guidelines.....	167, 173, 174
policy, routing	
omitting.....	66
usage guidelines.....	41
FTP traffic, sampling.....	323
G	
graceful Routing Engine switchover	
DHCP relay agent.....	304
group statement.....	375
usage guidelines.....	312
H	
hash-key statement.....	377
usage guidelines.....	281
hash-seed statement	
load balancing.....	408, 409
helpers statement.....	379
I	
icmp-code (firewall filter match condition).....	179
icmp-type (firewall filter match condition).....	176, 191
icons defined, notice.....	xxviii
if-exceeding statement.....	262
if-route-exists statement.....	147
import routing policies	
applying.....	26, 58, 59, 60, 151
overview.....	16
import statement.....	151
policy, routing	
usage guidelines.....	26, 58
indexed-next-hop statement.....	380
input statement.....	381
firewall filters	
usage guidelines.....	205, 225
forwarding table.....	381
port mirroring.....	381
sampling.....	382
traffic sampling.....	381
usage guidelines.....	288
usage guidelines.....	286
install-nexthop lsp (routing policy action).....	51
instance (routing policy match condition).....	43
instance-name.inet.0 routing table.....	18
interface (routing policy match condition).....	43
interface policers.....	245
interface set.....	245
interface statement	
accounting or sampling.....	383
BOOTP.....	384
DHCP relay agent.....	385
DNS or TFTP packet forwarding or relay	
agent.....	386
monitoring.....	387
next-hop group.....	387
port mirroring.....	388
usage guidelines.....	294, 312
interface-group (firewall filter match condition).....	176
interface-set statement.....	263
usage guidelines.....	245
interface-specific statement.....	263
usage guidelines.....	206
Internet Processor II ASIC	
determining presence of.....	277
invert-match statement.....	109
policy, routing	
usage guidelines.....	109
IP addresses	
sampling traffic from single IP addresses.....	322
ip-options (firewall filter match condition).....	190
ipsec-sa (firewall filter action).....	169
IPv6	
firewall filter match conditions.....	176, 179
IPv6 accounting, configuring.....	287
IS-IS	
policy, routing	
applying.....	27

J

joins, PIM
 rejecting.....122

L

layer2-unicast-replies statement.....388
 usage guidelines.....304
 LDP
 policy, routing
 applying.....27
 level (routing policy match condition).....43
 load balancing
 per-flow.....292, 293
 per-packet.....136
 IPv4.....137
 mpls.....138
 per-prefix.....293
 load-balance group.....248
 load-balance statement.....389
 usage guidelines.....293
 load-balance-group statement.....264
 usage guidelines.....248
 local-dump statement.....390
 usage guidelines.....327, 328
 local-preference
 policy, routing
 action.....51
 match condition.....43
 local-server-group statement.....391
 usage guidelines.....306
 log (firewall filter action).....171
 log output
 firewall filters.....172
 traffic sampling.....296, 326
 logical routers *See* logical systems
 logical systems
 configuring firewall filters.....208
 logical-bandwidth-policer statement.....264
 usage guidelines.....247
 logical-interface-policer statement.....265
 usage guidelines.....246, 251
 logical-system-name statement.....392
 longer route list match type.....113, 115
 loss-priority (firewall filter action).....169, 171

M

mac-address statement.....393
 manuals
 comments on.....xxvii
 match conditions
 firewall filters
 address filter.....184
 bit-field.....188, 189
 class-based filter.....190

 from statement.....167
 numeric range filter.....175
 overview.....160, 174
 policy, routing.....23, 41
 max-packets-per-second statement.....394
 maximum-capture-size statement.....393
 maximum-hop-count statement.....394
 usage guidelines.....294
 metric
 policy, routing
 action.....51
 match condition.....43
 minimum-wait-time statement.....395
 monitoring statement.....396
 usage guidelines.....290
 MPLS
 policy, routing
 applying.....27
 multicast-scoping
 policy, routing
 match condition.....44

N

names
 firewall filters.....166
 policy, routing.....40
 neighbor (routing policy match condition).....44
 next policy (routing policy control action).....48
 next term (routing policy control action).....48
 next-hop
 policy, routing
 action.....52
 match condition.....44
 next-hop groups.....291
 next-hop statement.....397
 next-hop groups
 usage guidelines.....291
 next-hop-group statement.....397
 usage guidelines.....291
 no-filter-check statement.....398
 no-label-1-exp statement
 usage guidelines.....140
 no-listen statement.....294
 usage guidelines.....295
 no-local-dump statement.....390, 398
 usage guidelines.....327
 no-stamp statement.....398, 423
 usage guidelines.....324
 no-world-readable statement.....399, 435
 usage guidelines.....326
 noncontiguous address filter.....187
 notice icons defined.....xxviii
 numeric range firewall filter match conditions.....175

O

option 82 information, DHCP	308
option-60 statement	399
option-82 statement	400
option 60 information, DHCP	306
origin	
policy, routing	
action	52
match condition	44
orlonger route list match type	113, 115
OSPF	
policy, routing	
applying	27
output files	
firewall log output file	172
logging information output file	296, 326
traffic sampling output files	324
output statement	401
accounting	401
firewall filters	
usage guidelines	205, 225
forwarding table	402
monitoring	402
port mirroring	403
sampling	404
usage guidelines	287
overrides statement	405
usage guidelines	304

P

packet capture	332
packet counter	
firewall filters	171
packet-capture statement	406
packet-length (firewall filter match condition)	176
packets	
testing packet protocols	191
parentheses, in syntax descriptions	xxix
password statement	407
per-flow load balancing	293
per-flow statement	408
per-prefix load balancing	293
per-prefix statement	409
PIM	
multicast traffic joins, rejecting	122
policy, routing	
applying	27
pipe ()	
bit-field logical operator	190
plus sign (+), bit-field logical operator	190
policer statement	
firewall	266
usage guidelines	236

policers	
example configurations	252
firewall filter action	169, 171
interface	245
policer action portion	236
policy (routing policy match condition)	45
policy framework	
architecture	6
comparison of policies	8
firewall filters	3
overview	3
policy, routing	3
policy, routing	
access and access-internal routes	57
actions	25, 46, 50, 56
applying	58, 150, 152
overview	25
architecture	6
AS path regular expressions	91, 97, 143, 144
BGP damping parameters	148
chains	
applying	59
evaluation	29
communities	97, 110, 145
comparison with firewall filters	8
configuring	36, 37, 39, 55, 56, 89
overview	22, 28
default policies and actions	19
evaluation	28
export policies	26, 149
flow, routing information	3
framework	15
import policies	16, 26, 64, 151
match conditions	23, 41, 46
multiple policies	
applying	59
evaluation	29
overview	15, 16
policy expressions	59, 64
preferences, modifying	52
prefix list	110, 155
prefix list filter	112, 155
purpose	8
rejecting PIM multicast traffic joins	122
route lists	114, 122
source prefixes, group	72
subroutines	30, 123, 126
terms	25, 40
testing	32
uses for	22
policy-options statement	152
usage guidelines	35
policy-statement statement	153
from statement	41
then statement	46, 47

- to statement.....41
- usage guidelines.....39
- port (firewall filter match condition).....176, 191
- port mirroring.....330
- port statement.....409
- usage guidelines.....327
- port-mirroring statement.....410
- usage guidelines.....331
- precedence (firewall filter match condition).....176, 179
- preferences
 - modifying
 - with routing policies.....52
 - policy, routing
 - action.....52
 - match condition.....45
- prefix list110, 155
- prefix list filter.....155
- prefix statement.....411
- usage guidelines.....308
- prefix-action statement.....267
- usage guidelines.....239
- prefix-length-range match type.....115
- prefix-list (routing policy match condition).....45
- prefix-list statement.....155
- usage guidelines.....186
- prefix-list-filter statement.....155
- protocols
 - applying policies.....25
 - firewall filter match condition.....179
 - match condition
 - firewall filters.....176
 - policy, routing.....45
 - routing
 - applying policies.....64, 66
 - testing packet protocols.....191

R

- rate statement.....412
- usage guidelines.....288
- reject
 - firewall filters
 - action.....171
 - policy, routing
 - control action.....48
- relay agents, DHCP
 - DHCP/BOOTP.....294
 - extended.....299
 - access and access-internal routes.....57, 303
 - active server groups.....311
 - clearing client address bindings and
 - statistics.....314
 - configuration examples.....317, 318, 319, 320
 - graceful Routing Engine switchover.....304
 - how components interact.....303
 - interface groups.....312

- option 82 information.....308
- option 60 information.....306
- overriding default configuration.....304
- server groups.....311
- state persistence.....304
- tracing operations.....314
- using external authentication.....313
- verifying client address bindings and
 - statistics.....314
- relay-option-60 statement.....413
- configuration example.....319, 320
- usage guidelines.....306
- relay-option-82 statement.....414
- usage guidelines.....308
- relay-server-group statement.....415
- usage guidelines.....306
- rib (routing policy match condition).....45
- RIP
 - policy, routing
 - applying.....27
- route lists.....114
- route recording.....326
- route-accounting statement.....416
- usage guidelines.....287
- route-filter (routing policy match condition).....46
- routers
 - DHCP relay agents.....294
- routing policy *See* policy, routing
- routing tables
 - instance-name.init.0.....18
- routing-instance (firewall filter action).....168, 171
- routing-instance statement.....416
- usage guidelines.....294
- routing-instance-name statement.....417
- RPF
 - firewall log and count.....173
- run-length statement.....417
- usage guidelines.....288

S

- sample (firewall filter action).....169, 171, 285
- sampled file.....296, 326
- sampled.pkts file.....324
- sampling statement.....418
- usage guidelines.....285
- server statement
 - DHCP and BOOTP service420
 - DNS and TFTP service420
 - usage guidelines.....294
- server-group statement.....421
- usage guidelines.....311
- service filters.....203
- service-filter statement
 - firewall.....268
- show chassis hardware command.....277

show firewall filter command.....	171
show interfaces policers command.....	245
show log command.....	172
show policers command.....	235
show policy damping command	133, 135
show route detail command	133
show route receive-protocol command	47
simple filters.....	203
simple-filter statement	
firewall.....	269
usage guidelines.....	203
single-rate statement	
usage guidelines.....	249
size statement	
packet capture.....	422
sampling.....	423
SONET interfaces	
sampling.....	321
source class usage.....	53, 72
source-address (firewall filter match	
condition).....	176, 179, 186
source-address-filter (routing policy match	
condition).....	46
source-class (routing policy action).....	53
source-port (firewall filter match	
condition).....	176, 179, 191
stamp option.....	325
stamp statement.....	423
usage guidelines.....	324
state persistence, DHCP.....	304
subroutines	30, 123
support, technical <i>See</i> technical support	
syntax conventions.....	xxviii
syslog (firewall filter action).....	169, 171

T

table statement.....	147
tag	
policy, routing	
action.....	53
tcp-established (firewall filter match condition).....	190
tcp-flags (firewall filter match condition).....	189
tcp-initial (firewall filter match condition).....	190
technical support	
contacting JTAC.....	xxxvii
term statement	
firewall.....	270
usage guidelines.....	166
policy	
usage guidelines.....	41
terms	
firewall filters.....	166
policy, routing.....	25, 40
test policy command	32

TFTP	
packet forwarding.....	295
requests, disabling recognition.....	295
tftp statement.....	424
then statement.....	153
firewall filters	
usage guidelines.....	168, 173
policy, routing	
usage guidelines.....	46, 47
usage guidelines.....	47
three-color-policer statement.....	273
usage guidelines.....	249
through route list match type.....	116
timestamp option.....	325
to statement.....	153
usage guidelines.....	41
topology (firewall filter action).....	168
trace (policy tracing action).....	47, 53
traceoptions statement	
DNS and TFTP packet forwarding.....	426
extended DHCP relay agent.....	428
usage guidelines.....	314
port mirroring and traffic sampling.....	430
usage guidelines.....	282
tracing actions.....	47, 53
tracing flags	
event policy	
all.....	298
configuration.....	298
database.....	298
events.....	298
policy.....	298
server.....	298
timer-events.....	298
traffic	
accounting.....	289
monitoring.....	290
sampling	
configuration statements.....	279, 335
disabling.....	321, 356
DNS and TFTP packet forwarding.....	295
example configurations.....	321
flow aggregation.....	326
FTP traffic.....	323
logging information output	
file.....	286, 296, 326
output files.....	324
overview.....	279
run-length parameter.....	288
sampling rate parameter.....	288
show log command.....	172
SONET interfaces.....	321
traffic from single IP addresses.....	322
traffic sampling	
configuring.....	285
traffic-class (firewall filter match condition).....	176, 179

tricolor marking policer.....	248
interface policer.....	251
trust-option-82 statement.....	430
usage guidelines.....	304
two-rate statement	
usage guidelines.....	249

U

unicast RPF.....	196
upto route list match type.....	116
user-prefix statement.....	432
username-include statement.....	431

V

var/log/jdhcpd file.....	314
var/log/sampled file.....	286, 296, 326
var/tmp/sampled.pkts file.....	324
vendor-option statement.....	433
usage guidelines.....	306
version statement.....	434
usage guidelines.....	327
version9 statement.....	435
usage guidelines.....	329
virtual-channel statement.....	274

W

world-readable statement.....	435
usage guidelines.....	326

Index of Statements and Commands

A

accounting statement.....	336
accounting-profile statement	
firewall.....	257
action statement	
policer (TCM).....	258
active-server-group statement.....	337
aggregation statement.....	338
always-write-giaddr statement.....	339
always-write-option-82 statement.....	340
apply-path statement.....	143
as-path statement.....	144
as-path-group statement.....	144
authentication statement.....	341
autonomous-system-type statement.....	342

B

bootp statement.....	343
----------------------	-----

C

cflowd statement.....	344
circuit-id statement.....	347
circuit-type statement.....	348
client-response-ttl statement.....	348
community statement.....	145
condition statement.....	147

D

damping statement.....	148
default-local-server-group statement.....	349
default-relay-server-group statement.....	350
delimiter statement.....	351
description statement	
interface.....	352
service.....	352
dhcp-relay statement.....	353
disable statement.....	356
disable-relay statement.....	357
domain statement.....	357
domain-name statement.....	358
drop statement.....	359

E

export statement.....	149
export-format statement.....	360

F

family inet statement (load balancing).....	363
family mpls statement.....	365
family multiservice statement.....	366
family statement.....	361
firewall filter.....	259
file statement	
helpers trace options.....	367
packet capture.....	368
sampling.....	368
traceoptions.....	369
filename statement.....	370
files statement	
packet capture.....	371
sampling.....	371
filter statement	
firewall.....	260
forwarding table.....	372
filter-specific statement.....	261
firewall statement.....	261
flood statement.....	372
flow-active-timeout statement.....	373
flow-export-destination statement.....	373
flow-inactive-timeout statement.....	374
forwarding-options statement.....	374
from statement.....	153

G

group statement.....	375
----------------------	-----

H

hash-key statement.....	377
helpers statement.....	379

I

if-exceeding statement.....	262
-----------------------------	-----

if-route-exists statement.....	147
import statement.....	151
indexed-next-hop statement.....	380
input statement.....	381
forwarding table.....	381
port mirroring.....	381
sampling.....	382
interface statement	
accounting or sampling.....	383
BOOTP.....	384
DHCP relay agent.....	385
DNS or TFTP packet forwarding or relay	
agent.....	386
monitoring.....	387
next-hop group.....	387
port mirroring.....	388
interface-set statement.....	263
interface-specific statement.....	263

L

layer2-unicast-replies statement.....	388
load-balance statement.....	389
load-balance-group statement.....	264
local-dump statement.....	390
local-server-group statement.....	391
logical-bandwidth-policer statement.....	264
logical-interface-policer statement.....	265
logical-system-name statement.....	392

M

mac-address statement.....	393
max-packets-per-second statement.....	394
maximum-capture-size statement.....	393
maximum-hop-count statement.....	394
minimum-wait-time statement.....	395
monitoring statement.....	396

N

next-hop statement.....	397
next-hop-group statement.....	397
no-filter-check statement.....	398
no-local-dump statement.....	390, 398
no-stamp statement.....	398, 423
no-world-readable statement.....	399, 435

O

option-60 statement.....	399
option-82 statement.....	400
output statement.....	401
accounting.....	401
forwarding table.....	402
monitoring.....	402

port mirroring.....	403
sampling.....	404
overrides statement.....	405

P

packet-capture statement.....	406
password statement.....	407
per-flow statement.....	408
per-prefix statement.....	409
policer statement	
firewall.....	266
policy-options statement.....	152
policy-statement statement.....	153
port statement.....	409
port-mirroring statement.....	410
prefix statement.....	411
prefix-action statement.....	267
prefix-list statement.....	155
prefix-list-filter statement.....	155

R

rate statement.....	412
relay-option-60 statement.....	413
relay-option-82 statement.....	414
relay-server-group statement.....	415
route-accounting statement.....	416
routing-instance statement.....	416
routing-instance-name statement.....	417
run-length statement.....	417

S

sampling statement.....	418
server statement	
DHCP and BOOTP service	420
DNS and TFTP service	420
server-group statement.....	421
service-filter statement	
firewall.....	268
show firewall filter command.....	171
show interfaces policers command.....	245
show log command.....	172
show policers command.....	235
show policy damping command	133, 135
show route detail command	133
show route receive-protocol command	47
simple-filter statement	
firewall.....	269
size statement	
packet capture.....	422
sampling.....	423

T

table statement.....	147
term statement	
firewall.....	270
test policy command	32
tftp statement.....	424
then statement.....	153
three-color-policer statement.....	273
to statement.....	153
traceoptions statement	
DNS and TFTP packet forwarding.....	426
extended DHCP relay agent.....	428
port mirroring and traffic sampling.....	430
trust-option-82 statement.....	430

U

user-prefix statement.....	432
username-include statement.....	431

V

vendor-option statement.....	433
version statement.....	434
version9 statement.....	435
virtual-channel statement.....	274

W

world-readable statement.....	435
-------------------------------	-----

