



---

Junos<sup>®</sup> OS

# Routing Policy Configuration Guide

Release

13.1



---

Published: 2013-02-10

Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

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

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

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® OS Routing Policy Configuration Guide*

13.1

Copyright © 2013, Juniper Networks, Inc.  
All rights reserved.

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

## YEAR 2000 NOTICE

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

## END USER LICENSE AGREEMENT

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

# Table of Contents

	About the Documentation . . . . .	xvii
	Documentation and Release Notes . . . . .	xvii
	Supported Platforms . . . . .	xvii
	Using the Examples in This Manual . . . . .	xviii
	Merging a Full Example . . . . .	xviii
	Merging a Snippet . . . . .	xix
	Documentation Conventions . . . . .	xix
	Documentation Feedback . . . . .	xxi
	Requesting Technical Support . . . . .	xxi
	Self-Help Online Tools and Resources . . . . .	xxi
	Opening a Case with JTAC . . . . .	xxii
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Introduction to Policy Framework . . . . .</b>	<b>3</b>
	Policy Framework Overview . . . . .	3
	Router Flows Affected by Policies . . . . .	4
	Policy Architecture . . . . .	6
	Control Points . . . . .	6
	Policy Components . . . . .	7
	Default Policies and Actions . . . . .	8
	Routing Policies Configuration Overview . . . . .	10
	Policy Configuration Recommendations . . . . .	11
	Comparison of Routing Policies and Firewall Filters . . . . .	11
<b>Chapter 2</b>	<b>Introduction to Routing Policy . . . . .</b>	<b>15</b>
	Understanding Routing Policies . . . . .	15
	Importing and Exporting Routes . . . . .	17
	Default Routing Policies and Actions . . . . .	18
	Creating Routing Policies . . . . .	20
	Configuring a Routing Policy . . . . .	20
	Routing Policy Named Match Conditions . . . . .	21
	Route-Based Match Conditions . . . . .	21
	Protocol-Based Match Conditions . . . . .	22
	Routing Policy Actions . . . . .	23
	Routing Policy Terms . . . . .	23
	Applying a Routing Policy . . . . .	24
	Routing Protocol Support for Import and Export Policies . . . . .	24
	Applying a Routing Policy to Routing Protocols . . . . .	25
	Applying Export Policies to the Forwarding Table . . . . .	25
	Evaluating a Routing Policy . . . . .	26

	How a Routing Policy Is Evaluated . . . . .	26
	How a Routing Policy Chain Is Evaluated . . . . .	27
	How a Routing Policy Expression Is Evaluated . . . . .	28
	How a Routing Policy Subroutine Is Evaluated . . . . .	28
	Routing Policy Tests . . . . .	30
<b>Chapter 3</b>	<b>Introduction to Extended Match Conditions Configuration . . . . .</b>	<b>31</b>
	Understanding BGP Communities and Extended Communities as Routing Policy	
	Match Conditions . . . . .	31
	How BGP Communities and Extended Communities Are Evaluated in Routing	
	Policy Match Conditions . . . . .	32
	Multiple Matches . . . . .	33
	Extended Community Type . . . . .	34
	Multiple Communities Are Matched with Ex-OR Logic . . . . .	35
<b>Chapter 4</b>	<b>Introduction to Extended Actions Configuration . . . . .</b>	<b>37</b>
	Overview of Per-Packet Load Balancing . . . . .	37
<b>Chapter 5</b>	<b>Introduction to Traffic Sampling, Forwarding, and Monitoring Overview . . . . .</b>	<b>39</b>
	Traffic Sampling, Forwarding, and Monitoring Overview . . . . .	39
<b>Chapter 6</b>	<b>Introduction to Traffic Forwarding and Monitoring Configuration . . . . .</b>	<b>41</b>
	Per-Flow and Per-Prefix Load Balancing Overview . . . . .	41
<b>Part 2</b>	<b>Configuration</b>	
<b>Chapter 7</b>	<b>Routing Policy Configuration Statements . . . . .</b>	<b>45</b>
	Configuring a Routing Policy . . . . .	45
	Minimum Routing Policy Configuration . . . . .	46
	Minimum Routing Policy Chain Configuration . . . . .	46
	Minimum Subroutine Configuration . . . . .	47
<b>Chapter 8</b>	<b>Routing Policy Configuration . . . . .</b>	<b>49</b>
	Defining Routing Policies . . . . .	50
	Configuring Match Conditions in Routing Policy Terms . . . . .	51
	Configuring Actions in Routing Policy Terms . . . . .	58
	Configuring Flow Control Actions . . . . .	60
	Configuring Actions That Manipulate Route Characteristics . . . . .	60
	Configuring the Default Action in Routing Policies . . . . .	67
	Example: Configuring the Default Action in a Routing Policy . . . . .	68
	Configuring a Final Action in Routing Policies . . . . .	68
	Logging Matches to a Routing Policy Term . . . . .	69
	Configuring Separate Actions for Routes in Route Lists . . . . .	69
	Applying Routing Policies and Policy Chains to Routing Protocols . . . . .	70
	Applying Policy Expressions to Routes Exported from Routing Tables . . . . .	70
	Policy Expression Examples . . . . .	72
	How a Policy Expression Is Evaluated . . . . .	73
	Example: Evaluating Policy Expressions . . . . .	74
	Applying Routing Policies to the Forwarding Table . . . . .	75

Configuring Dynamic Routing Policies . . . . .	76
Configuring Routing Policies and Policy Objects in the Dynamic Database . . . . .	77
Configuring Routing Policies Based on Dynamic Database Configuration . . . . .	77
Applying Dynamic Routing Policies to BGP . . . . .	79
Preventing Reestablishment of BGP Peering Sessions After NSR Routing Engine Switchover . . . . .	79
Example: Configuring a BGP Export Policy That References a Dynamic Routing Policy . . . . .	80
Forwarding Packets to the Discard Interface . . . . .	81
Testing Routing Policies . . . . .	82
Example: Testing a Routing Policy . . . . .	83
Example: Creating a Routing Policy . . . . .	83
Routing Policy Examples . . . . .	84
Example: Creating a Routing Policy Term . . . . .	85
Example: Rejecting Known Invalid Routes . . . . .	86
Examples: Configuring Static Routes . . . . .	87
Understanding Basic Static Routing . . . . .	88
Example: Configuring a Basic Set of Static Routes . . . . .	88
Example: Configuring IPv6 Static Routes . . . . .	93
Example: Defining a Routing Policy from BGP to IS-IS . . . . .	97
Example: Using Routing Policy to Set a Preference . . . . .	98
Example: Importing and Exporting Access and Access-Internal Routes in a Routing Policy . . . . .	98
Example: Exporting Routes to IS-IS . . . . .	99
Example: Applying Export and Import Policies to BGP Peer Groups . . . . .	99
Example: Defining a Routing Policy Based on the Number of BGP Communities . . . . .	100
Example: Applying a Prefix to Routes Learned from a Peer . . . . .	100
Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS . . . . .	100
Example: Redistributing OSPF Routes into BGP . . . . .	101
Example: Exporting Direct Routes Into IS-IS . . . . .	102
Example: Exporting Internal IS-IS Level 1 Routes to Level 2 . . . . .	102
Example: Exporting IS-IS Level 2 Routes to Level 1 . . . . .	102
Example: Assigning Different Forwarding Next-Hop LSPs to Different Destination Prefixes . . . . .	103
Example: Grouping Destination Prefixes . . . . .	104
Example: Grouping Source Prefixes . . . . .	105
Example: Grouping Source and Destination Prefixes into a Forwarding Class . . . . .	106
Example: Accepting Routes with Specific Destination Prefixes . . . . .	108
Example: Accepting Routes from BGP with a Specific Destination Prefix . . . . .	108
Example: Using Routing Policy in an ISP Network . . . . .	109
Requesting a Single Default Route on the Customer 1 Router . . . . .	111
Requesting Specific Routes on the Customer 2 Router . . . . .	112
Configuring a Peer Policy on ISP Router 3 . . . . .	114
Configuring Private and Exchange Peers on ISP Router 1 and 2 . . . . .	116
Configuring Locally Defined Static Routes on the Exchange Peer 2 Router . . . . .	119

**Chapter 9**

Configuring Outbound and Generated Routes on the Private Peer 2 Router	119
Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Switches	121
Understanding the Default BGP Routing Policy on Packet Transport Switches	121
Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Switches	123
Understanding ECMP Flow-Based Forwarding	126
Example: Configuring ECMP Flow-Based Forwarding	127
<b>Extended Match Conditions Configuration</b>	<b>133</b>
Configuring AS Path Regular Expressions to Use as Routing Policy Match Conditions	133
Configuring AS Path Regular Expressions	134
Configuring a Null AS Path	138
How AS Path Regular Expressions Are Evaluated	139
Examples: Configuring AS Path Regular Expressions	139
Defining BGP Communities and Extended Communities for Use in Routing Policy Match Conditions	140
Defining BGP Communities for Use in Routing Policy Match Conditions	141
Using UNIX Regular Expressions in Community Names	142
Defining BGP Extended Communities for Use in Routing Policy Match Conditions	144
Examples: Defining BGP Extended Communities	146
Inverting Community Matches	146
Including BGP Communities and Extended Communities in Routing Policy Match Conditions	146
Using Routing Policies to Prevent Advertisement of BGP Communities to Neighbors	147
Examples: Configuring BGP Communities as Routing Policy Match Conditions	147
Configuring Prefix Lists for Use in Routing Policy Match Conditions	151
Configuring Prefix Lists	152
How Prefix Lists Are Evaluated in Routing Policy Match Conditions	154
Configuring Prefix List Filters	154
Example: Configuring a Prefix List	155
Configuring Route Lists for Use in Routing Policy Match Conditions	155
Configuring Route Lists	156
How Route Lists Are Evaluated in Routing Policy Match Conditions	160
How an Address Mask Match Type Is Evaluated	161
How Prefix Order Affects Route List Evaluation	162
Common Configuration Problem with the Longest-Match Lookup	162
Route List Examples	163
Example: Rejecting Routes with Specific Destination Prefixes and Mask Lengths	163
Example: Rejecting Routes with a Mask Length Greater than Eight	164
Example: Rejecting Routes with Mask Length Between 26 and 29	164
Example: Rejecting Routes from Specific Hosts	164

	Example: Accepting Routes with a Defined Set of Prefixes . . . . .	165
	Example: Rejecting Routes with a Defined Set of Prefixes . . . . .	165
	Example: Rejecting Routes with Prefixes Longer than 24 Bits . . . . .	166
	Example: Rejecting PIM Multicast Traffic Joins . . . . .	166
	Example: Rejecting PIM Traffic . . . . .	166
	Example: Accepting Incoming IPv4 Routes by Applying an Address Mask to the Route Address and the Destination Match Prefix . . . . .	167
	Example: Accepting Incoming IPv4 Routes with Similar Patterns But Different Prefix Lengths . . . . .	168
	Example: Evaluation of an Address Mask Match Type with Longest-Match Lookup . . . . .	169
	Configuring Subroutines in Routing Policy Match Conditions . . . . .	170
	Configuring Subroutines . . . . .	170
	Possible Consequences of Termination Actions in Subroutines . . . . .	171
	Example: Configuring a Subroutine . . . . .	173
	Configuring Routing Policy Match Conditions Based on Routing Table Entries . .	174
<b>Chapter 10</b>	<b>Extended Actions Configuration . . . . .</b>	<b>177</b>
	Prepending AS Numbers to BGP AS Paths . . . . .	177
	Example: Configuring a Routing Policy to Prepend the AS Path . . . . .	178
	Adding AS Numbers to BGP AS Paths . . . . .	180
	Using Routing Policies to Damp BGP Route Flapping . . . . .	181
	Configuring BGP Flap Damping Parameters . . . . .	181
	Specifying BGP Flap Damping as the Action in Routing Policy Terms . . . .	184
	Disabling Damping for Specific Address Prefixes . . . . .	184
	Example: Disabling Damping for a Specific Address Prefix . . . . .	185
	Example: Configuring BGP Flap Damping . . . . .	185
	Configuring Per-Packet Load Balancing . . . . .	186
	Per-Packet Load Balancing Examples . . . . .	188
	Configuring Load Balancing Based on MPLS Labels on DPC I-Chip-Based Hardware . . . . .	189
	Configuring Load Balancing for Ethernet Pseudowires . . . . .	193
	Configuring Load Balancing Based on MAC Addresses . . . . .	194
	Configuring VPLS Load Balancing Based on IP and MPLS Information . . . . .	195
	Configuring VPLS Load Balancing on MX Series 3D Universal Edge Routers . .	196
	Example: Configuring Multicast Load Balancing over Aggregated Ethernet Links . . . . .	197
<b>Chapter 11</b>	<b>Traffic Sampling Configuration . . . . .</b>	<b>211</b>
	Traffic Sampling Configuration . . . . .	211
	Minimum Traffic Sampling Configuration . . . . .	212
	Configuring Traffic Sampling . . . . .	213
	Disabling Traffic Sampling . . . . .	215
	Configuring the Output File for Traffic Sampling . . . . .	215
	Traffic Sampling Output Format . . . . .	216
	Tracing Traffic-Sampling Operations . . . . .	217
	Configuring Flow Aggregation (cflowd) . . . . .	217
	Debugging cflowd Flow Aggregation . . . . .	219

	Configuring Active Flow Monitoring Using Version 9 . . . . .	220
	Example: Configuring Active Flow Monitoring Using Version 9 . . . . .	221
	Traffic Sampling Examples . . . . .	221
	Example: Sampling a Single SONET/SDH Interface . . . . .	221
	Example: Sampling All Traffic from a Single IP Address . . . . .	222
	Example: Sampling All FTP Traffic . . . . .	223
<b>Chapter 12</b>	<b>Traffic Forwarding and Monitoring Configuration . . . . .</b>	<b>225</b>
	Configuring Traffic Forwarding and Monitoring . . . . .	225
	Configuring Forwarding Table Filters . . . . .	229
	Overview of Forwarding Table Filters . . . . .	229
	Configuring a Forwarding Table Filter . . . . .	230
	Applying Filters to Forwarding Tables . . . . .	231
	Configuring IPv6 Accounting . . . . .	234
	Configuring Discard Accounting . . . . .	234
	Configuring Flow Monitoring . . . . .	235
	Forwarding Table Filters for Routing Instances on ACX Series Routers . . . . .	236
	Configuring Load-Balance Groups . . . . .	237
	Configuring Next-Hop Groups . . . . .	237
	Configuring Per-Prefix Load Balancing . . . . .	238
	Configuring Per-Flow Load Balancing Based on Hash Values . . . . .	239
	Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents . . . . .	239
	Configuring DNS and TFTP Packet Forwarding . . . . .	241
	Tracing BOOTP, DNS, and TFTP Forwarding Operations . . . . .	242
	Configuring the Log Filename . . . . .	243
	Configuring the Number and Size of Log Files . . . . .	243
	Configuring Access to the Log File . . . . .	244
	Configuring a Regular Expression for Lines to Be Logged . . . . .	244
	Example: Configuring DNS Packet Forwarding . . . . .	244
	Preventing DHCP Spoofing on MX Series 3D Universal Edge Routers . . . . .	245
	Configuring Port Mirroring . . . . .	246
	Configuration Guidelines . . . . .	247
	Configuring Port Mirroring . . . . .	248
	Configuring the Port-Mirroring Address Family and Interface . . . . .	248
	Configuring Multiple Port-Mirroring Instances . . . . .	248
	Configuring Port-Mirroring Instances . . . . .	249
	Associating a Port-Mirroring Instance on M320 Routers . . . . .	249
	Associating a Port-Mirroring Instance on M120 Routers . . . . .	249
	Configuring MX Series 3D Universal Edge Routers and M120 Routers to Mirror Traffic Only Once . . . . .	250
	Configuring Packet Capture . . . . .	250
<b>Part 3</b>	<b>Administration</b>	
<b>Chapter 13</b>	<b>Routing Policy Reference . . . . .</b>	<b>255</b>
	Protocols That Can Be Imported to and Exported from the Routing Table . . . . .	255
	Routing Tables Affected by Routing Policies . . . . .	256
	Default Import and Export Policies for Protocols . . . . .	257
	Routing Policy Match Conditions . . . . .	258



	Protocol Support for Import and Export Policies . . . . .	261
<b>Chapter 14</b>	<b>Summary of Routing Policy Configuration Statements . . . . .</b>	<b>263</b>
	aigp-originate . . . . .	264
	apply-path . . . . .	265
	as-path (Policy Options) . . . . .	266
	as-path-group . . . . .	267
	community (Policy Options) . . . . .	268
	condition . . . . .	270
	damping (Policy Options) . . . . .	271
	dynamic-db . . . . .	272
	export (Protocols BGP) . . . . .	273
	export (Protocols IS-IS) . . . . .	274
	export (Protocols DVMRP) . . . . .	275
	export (Protocols LDP) . . . . .	275
	export (Protocols MSDP) . . . . .	276
	export (Protocols OSPF) . . . . .	277
	export (Protocols PIM) . . . . .	278
	export (Protocols PIM Bootstrap) . . . . .	278
	export (Protocols RIP) . . . . .	279
	export (Protocols RIPng) . . . . .	280
	export (Routing Options) . . . . .	281
	import (Protocols BGP) . . . . .	282
	import (Protocols DVMRP) . . . . .	283
	import (Protocols LDP) . . . . .	283
	import (Protocols MSDP) . . . . .	284
	import (Protocols OSPF) . . . . .	285
	import (Protocols PIM) . . . . .	286
	import (Protocols PIM Bootstrap) . . . . .	287
	import (Protocols RIP) . . . . .	288
	import (Protocols RIPng) . . . . .	289
	import (Routing Options) . . . . .	290
	policy-options . . . . .	290
	policy-statement . . . . .	291
	prefix-list . . . . .	293
	prefix-list-filter . . . . .	294
	rtf-prefix-list . . . . .	295
<b>Chapter 15</b>	<b>Summary of Traffic Sampling, Forwarding, and Monitoring Configuration Statements . . . . .</b>	<b>297</b>
	accounting . . . . .	298
	aggregation . . . . .	299
	autonomous-system-type . . . . .	300
	bootp . . . . .	301
	cflowd (Discard Accounting) . . . . .	302
	cflowd (Flow Monitoring) . . . . .	303
	client-response-ttl . . . . .	303
	description (Forwarding Options) . . . . .	304
	dhcp-relay (DHCP Spoofing Prevention) . . . . .	305
	disable (Forwarding Options) . . . . .	305

domain	306
export-format	306
enhanced-hash-key	307
family (Filtering)	310
family (Monitoring)	311
family (Port Mirroring)	312
family (Sampling)	313
family inet	315
family mpls	317
family multiservice	320
file (Extended DHCP Relay Agent and Helpers Trace Options)	322
file (Packet Capture)	322
file (Sampling)	323
file (Trace Options)	323
filename (Packet Capture)	324
filename (Sampling)	324
files (Packet Capture)	325
files (Sampling and Traceoptions)	325
filter (IPv4, IPv6, and MPLS)	326
filter (VPLS)	326
flood	327
flow-active-timeout	327
flow-export-destination	328
flow-inactive-timeout	328
flow-server	329
forwarding-options	330
group (DHCP Spoofing Prevention)	330
hash-key (Forwarding Options)	331
helpers	333
indexed-load-balance	335
input (Forwarding Table)	336
input (Port Mirroring)	336
input (Sampling)	337
instance	338
interface (Accounting or Sampling)	339
interface (BOOTP)	340
interface (DHCP Spoofing Prevention)	341
interface (DNS and TFTP Packet Forwarding or Relay Agent)	342
interface (Monitoring)	343
interface (Next-Hop Group)	344
interface (Port Mirroring)	344
load-balance (Forwarding Options)	345
load-balance-group	347
local-dump	347
max-packets-per-second	348
maximum-capture-size	348
maximum-hop-count	349
maximum-packet-length	350
minimum-wait-time	351

mirror-once .....	351
monitoring .....	352
next-hop (Forwarding Options) .....	353
next-hop-group .....	354
no-filter-check .....	355
no-listen .....	355
output (Accounting) .....	356
output (Forwarding Table) .....	357
output (Monitoring) .....	358
output (Port Mirroring) .....	359
output (Sampling) .....	360
packet-capture .....	361
per-flow .....	362
per-prefix .....	363
port (cflowd) .....	363
port (Packet Forwarding) .....	364
port-mirroring .....	366
rate (Forwarding Options) .....	367
relay-agent-option .....	368
route-accounting .....	368
run-length .....	369
sampling (Forwarding Options) .....	370
server (DHCP and BOOTP Relay Agent) .....	372
server (DNS and TFTP Service) .....	373
size (Packet Capture) .....	373
size (Sampling and Traceoptions) .....	374
stamp .....	375
tftp .....	375
traceoptions (DNS and TFTP Packet Forwarding) .....	376
traceoptions (Port Mirroring and Traffic Sampling) .....	378
version .....	378
version9 .....	379
world-readable (Forwarding Options) .....	380

## Part 4

## Index

Index .....	383
-------------	-----



# List of Figures

<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Introduction to Policy Framework</b>	<b>3</b>
	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	6
	Figure 4: Policy Control Points	7
<b>Chapter 2</b>	<b>Introduction to Routing Policy</b>	<b>15</b>
	Figure 5: Importing and Exporting Routing Policies	16
	Figure 6: Importing and Exporting Routes	17
	Figure 7: Importing and Exporting Routing Policies	20
	Figure 8: Routing Policy Components	21
	Figure 9: Routing Policy Evaluation	27
	Figure 10: Routing Policy Chain Evaluation	28
	Figure 11: Routing Policy Subroutine Evaluation	30
<b>Part 2</b>	<b>Configuration</b>	
<b>Chapter 8</b>	<b>Routing Policy Configuration</b>	<b>49</b>
	Figure 12: Customer Routes Connected to a Service Provider	89
	Figure 13: Customer Routes Connected to a Service Provider	94
	Figure 14: ISP Network Example	110
	Figure 15: ECMP Routes	128
<b>Chapter 10</b>	<b>Extended Actions Configuration</b>	<b>177</b>
	Figure 16: Multicast Load Balancing over Aggregated Ethernet Links	198
<b>Chapter 11</b>	<b>Traffic Sampling Configuration</b>	<b>211</b>
	Figure 17: Configure Sampling Rate	214



# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xvii</b>
	Table 1: Notice Icons . . . . .	xix
	Table 2: Text and Syntax Conventions . . . . .	xx
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Introduction to Policy Framework</b> . . . . .	<b>3</b>
	Table 3: Summary of Key Routing Policy Actions . . . . .	9
	Table 4: Purpose of Routing Policies and Firewall Filters . . . . .	11
	Table 5: Implementation Differences Between Routing Policies and Firewall Filters . . . . .	12
<b>Chapter 2</b>	<b>Introduction to Routing Policy</b> . . . . .	<b>15</b>
	Table 6: Route List Match Types . . . . .	22
<b>Part 2</b>	<b>Configuration</b>	
<b>Chapter 8</b>	<b>Routing Policy Configuration</b> . . . . .	<b>49</b>
	Table 7: Routing Policy Match Conditions . . . . .	52
	Table 8: Flow Control Actions . . . . .	60
	Table 9: Actions That Manipulate Route Characteristics . . . . .	61
	Table 10: Policy Action Conversion Values . . . . .	71
	Table 11: Policy Expression Logical Operators . . . . .	71
<b>Chapter 9</b>	<b>Extended Match Conditions Configuration</b> . . . . .	<b>133</b>
	Table 12: AS Path Regular Expression Operators . . . . .	136
	Table 13: Examples of AS Path Regular Expressions . . . . .	136
	Table 14: Community Attribute Regular Expression Operators . . . . .	143
	Table 15: Examples of Community Attribute Regular Expressions . . . . .	144
	Table 16: Prefix List and Route List Differences . . . . .	152
	Table 17: Route List Match Types for a Prefix List Filter . . . . .	154
	Table 18: Route List Match Types for a Prefix List . . . . .	158
	Table 19: Match Type Examples . . . . .	159
<b>Chapter 10</b>	<b>Extended Actions Configuration</b> . . . . .	<b>177</b>
	Table 20: Damping Parameters . . . . .	182
	Table 21: MPLS LSP Load Balancing Options . . . . .	190
<b>Part 3</b>	<b>Administration</b>	
<b>Chapter 13</b>	<b>Routing Policy Reference</b> . . . . .	<b>255</b>

Table 22: Protocols That Can Be Imported to and Exported from the Routing Table . . . . .	255
Table 23: Routing Tables Affected by Routing Policies . . . . .	256
Table 24: Default Import and Export Policies for Protocols . . . . .	257
Table 25: Match Conditions . . . . .	258
Table 26: Summary of Key Routing Policy Match Conditions . . . . .	260
Table 27: Protocol Support for Import and Export Policies . . . . .	261



# About the Documentation

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

## Documentation and Release Notes

---

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

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

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

## Supported Platforms

---

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

- M Series
- T Series
- MX Series
- J Series
- SRX Series
- PTX Series

## Using the Examples in This Manual

---

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

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

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

### Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

## Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

## Documentation Conventions

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

Table 1: Notice Icons





Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.

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

Table 2: Text and Syntax Conventions

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

## J-Web GUI Conventions

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

Convention	Description	Examples
<b>Bold text like this</b>	Represents J-Web graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> <li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li> <li>To cancel the configuration, click <b>Cancel</b>.</li> </ul>
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## 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](mailto:techpubs-comments@juniper.net), or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

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

## Requesting Technical Support

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

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

## Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>

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

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

## Opening a Case with JTAC

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

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

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

## PART 1

# Overview

- [Introduction to Policy Framework on page 3](#)
- [Introduction to Routing Policy on page 15](#)
- [Introduction to Extended Match Conditions Configuration on page 31](#)
- [Introduction to Extended Actions Configuration on page 37](#)
- [Introduction to Traffic Sampling, Forwarding, and Monitoring Overview on page 39](#)
- [Introduction to Traffic Forwarding and Monitoring Configuration on page 41](#)





## CHAPTER 1

# Introduction to Policy Framework

- Policy Framework Overview on page 3
- Router Flows Affected by Policies on page 4
- Policy Architecture on page 6
- Control Points on page 6
- Policy Components on page 7
- Default Policies and Actions on page 8
- Routing Policies Configuration Overview on page 10
- Policy Configuration Recommendations on page 11
- Comparison of Routing Policies and Firewall Filters on page 11

## Policy Framework Overview

---

The Junos<sup>®</sup> operating system (Junos OS) provides a *policy framework*, which is a collection of Junos OS 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.

## Router Flows Affected by Policies

The Junos OS 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**




1630

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.)

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 BGP route flap-damping values.
- Perform per-packet load balancing.
- Enable class of service (CoS).

**Figure 2: Routing Policies to Control Routing Information Flow**

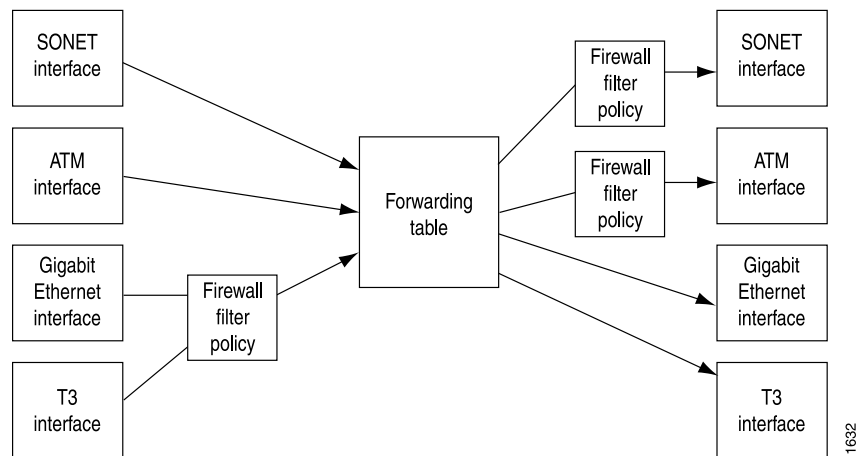


You can configure firewall filters to control the following aspects of packet flow (see [Figure 3 on page 6](#)):

- 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 OS 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 OS 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 8](#)
- [Routing Policies Configuration Overview on page 10](#)
- [Policy Configuration Recommendations on page 11](#)



**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 11](#).

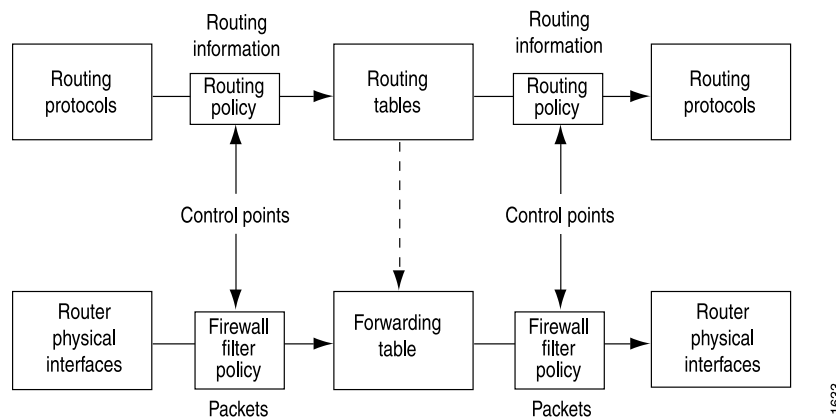
## Control Points

All policies provide two points at which you can control routing information or packets through the router (see [Figure 4 on page 7](#)). 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 7 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 20 and Stateless Firewall Filter Components.

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 11.](#)) 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.

An action defines what the router does with the route when the route matches all the match conditions in the **from** and **to** statements for a particular term. If a term does not have **from** and **to** statements, all routes are considered to match and the actions apply to all routes.

Each term can have one or more of the following types of actions. The actions are configured under the **then** statement.

- 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

If you do not specify an action, one of the following results occurs:

- The next term in the routing policy, if one exists, is evaluated.
- If the routing policy has no more terms, the next routing policy, if one exists, is evaluated.
- If there are no more terms or routing policies, the accept or reject action specified by the default policy is executed.

[Table 3 on page 9](#) summarizes the routing policy actions.

Table 3: Summary of Key Routing Policy Actions

Action	Description
<b>Flow Control Actions</b>	These actions control the flow of routing information into and out of the routing table.
<b>accept</b>	Accepts the route and propagates it. After a route is accepted, no other terms in the routing policy and no other routing policies are evaluated.
<b>reject</b>	Rejects the route and does not propagate it. After a route is rejected, no other terms in the routing policy and no other routing policies are evaluated.
<b>next term</b>	Skips to and evaluates the next term in the same routing policy. Any <b>accept</b> or <b>reject</b> action specified in the <b>then</b> statement is ignored. Any actions specified in the <b>then</b> statement that manipulate route characteristics are applied to the route.
<b>next policy</b>	Skips to and evaluates the next routing policy. Any <b>accept</b> or <b>reject</b> action specified in the <b>then</b> statement is ignored. Any actions specified in the <b>then</b> statement that manipulate route characteristics are applied to the route.
<b>Route Manipulation Actions</b>	These actions manipulate the route characteristics.
<b>as-path-prepend <i>as-path</i></b>	<p>Appends one or more AS numbers at the beginning of the AS path. If you are specifying more than one AS number, include the numbers in quotation marks.</p> <p>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 appended AS numbers are placed within a confederation sequence. Otherwise, the appended AS numbers are placed with a nonconfederation sequence.</p>
<b>as-path-expand last-as count <i>n</i></b>	<p>Extracts the last AS number in the existing AS path and appends that AS number to the beginning of the AS path <i>n</i> times. Replace <i>n</i> with a number from 1 through 32.</p> <p>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 appended AS numbers are placed within a confederation sequence. Otherwise, the appended AS numbers are placed with a nonconfederation sequence.</p>
<b>class <i>class-name</i></b>	Applies the specified class-of-service (CoS) parameters to routes installed into the routing table.
<b>color <i>preference</i></b> <b>color2 <i>preference</i></b>	Sets the preference value to the specified value. The <b>color</b> and <b>color2</b> preference values can be a number from 0 through 4,294,967,295 ( $2^{32} - 1$ ). A lower number indicates a more preferred route.
<b>damping <i>name</i></b>	<p>Applies the specified route-damping parameters to the route. These parameters override BGP's default damping parameters.</p> <p>This action is useful only in import policies.</p>
<b>local-preference <i>value</i></b>	Sets the BGP local preference attribute. The preference can be a number from 0 through 4,294,967,295 ( $2^{32} - 1$ ).

Table 3: Summary of Key Routing Policy Actions (*continued*)

Action	Description
<b>metric</b> <i>metric</i>	Sets the metric. You can specify up to four metric values, starting with <b>metric</b> (for the first metric value) and continuing with <b>metric2</b> , <b>metric3</b> , and <b>metric4</b> .
<b>metric2</b> <i>metric</i>	
<b>metric3</b> <i>metric</i>	For BGP routes, <b>metric</b> corresponds to the MED, and <b>metric2</b> corresponds to the IGP metric if the BGP next hop loops through another router.
<b>metric4</b> <i>metric</i>	
<b>next-hop</b> <i>address</i>	Sets the next hop.  If you specify <b>address</b> as <b>self</b> , the next-hop address is replaced by one of the local router's addresses. The advertising protocol determines which address to use.

## Routing Policies Configuration Overview

To configure a routing policy:

1. Determine what you want to accomplish with the policy, and thoroughly understand how to achieve your goal using the various match conditions and actions. Define the policy and 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](#).
2. Make certain that you understand the default policies and actions for the policy you are configuring.



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

3. Apply the policy to whatever moves the routing information or packets through the router, for example, the routing protocol or the router interface.
4. Configure an interface on the router. See the *Junos OS Interfaces Configuration Guide for Security Devices*.
5. Configure an Interior Gateway Protocol (IGP) and Border Gateway Protocol (BGP), if necessary. See:
  - RIP Configuration Overview
  - OSPF Configuration Overview
  - Example: Configuring IS-IS
  - BGP Configuration Overview
6. Configure the router interface to reject or accept routes, if necessary.
7. Configure static routes, if necessary. See [“Example: Configuring a Basic Set of Static Routes” on page 88](#).
8. Name the policy. See [“Example: Creating a Routing Policy” on page 83](#).



9. Configure the policy term. See [“Example: Creating a Routing Policy Term” on page 85](#).
10. (Optional) Reject useless routes. See [“Example: Rejecting Known Invalid Routes” on page 86](#).
11. (Optional) Advertise additional routes. See [Example: Injecting OSPF Routes into the BGP Routing Table](#).
12. (Optional) Create a forwarding class. See [“Example: Grouping Source and Destination Prefixes into a Forwarding Class” on page 106](#).
13. (Optional) Make a route less preferable to BGP. See [“Example: Configuring a Routing Policy to Prepend the AS Path” on page 178](#).
14. (Optional) Suppress route information. See [Example: Configuring Damping Parameters](#).

#### Related Documentation

- [Understanding Routing Policies on page 15](#)
- [Minimum Routing Policy Configuration on page 46](#)
- [Testing Routing Policies on page 82](#)

## Policy Configuration Recommendations

The Junos OS 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 4 on page 11](#) describes their purposes. [Table 5 on page 12](#) compares the implementation details for routing policies and firewall filters, highlighting the similarities and differences in their configuration.

**Table 4: 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.

Table 4: Purpose of Routing Policies and Firewall Filters (*continued*)

Policies	Source	Policy Purpose
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 5: 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: <ul style="list-style-type: none"> <li>Define policy</li> <li>Apply policy</li> </ul>	<p>Define a policy that contains terms, match conditions, and actions.</p> <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>Define a policy that contains terms, match conditions, and actions.</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 5: 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"> <li>• Autonomous system (AS) path expression—A combination of AS numbers and regular expression operators.</li> <li>• Community—A group of destinations that share a common property.</li> <li>• Prefix list—A named list of prefixes.</li> <li>• Route list—A list of destination prefixes.</li> <li>• Subroutine—A routing policy that is called repeatedly from other routing policies.</li> </ul>	<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"> <li>• Numeric values, such as port and protocol numbers.</li> <li>• Prefix values, such as IP source and destination prefixes.</li> <li>• 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.</li> </ul>
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"> <li>• Accept—Accept the route into the routing table, and propagate it. After this action is taken, the evaluation of subsequent terms and policies ends.</li> <li>• 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.</li> </ul> <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"> <li>• Next term—Evaluate the next term in the routing policy.</li> <li>• Next policy—Evaluate the next routing policy.</li> <li>• 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.</li> <li>• Trace action, which logs route matches.</li> </ul>	<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"> <li>• Accept—Accept a packet.</li> <li>• Discard—Discard a packet silently, without sending an ICMP message.</li> <li>• Reject—Discard a packet, and send an ICMP destination unreachable message.</li> <li>• Routing instance—Specify a routing table to which packets are forwarded.</li> <li>• Next term—Evaluate the next term in the firewall filter.</li> </ul> <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"> <li>• Count—Add packet to a count total.</li> <li>• Forwarding class—Set the packet forwarding class to a specified value from 0 through 3.</li> <li>• IPsec security association—Used with the source and destination address match conditions, specify an IP Security (IPsec) security association (SA) for the packet.</li> <li>• Log—Store the header information of a packet on the Routing Engine.</li> <li>• Loss priority—Set the packet loss priority (PLP) bit to a specified value, 0 or 1.</li> <li>• Policer—Apply rate-limiting procedures to the traffic.</li> <li>• Sample—Sample the packet traffic.</li> <li>• Syslog—Log an alert for the packet.</li> </ul>

**Table 5: 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"> <li>• If a policy does not specify a match condition, all routes evaluated against the policy match.</li> <li>• 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"> <li>• The next term, if present, is evaluated.</li> <li>• If no other terms are present, the next policy is evaluated.</li> <li>• If no other policies are present, the action specified by the default policy is taken.</li> </ul> </li> <li>• 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.</li> <li>• If a match does not occur with any terms in a policy and subsequent policies exist, the next policy is evaluated.</li> <li>• 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.</li> </ul>	<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"> <li>• If a firewall filter does not specify a match condition, all packets are considered to match.</li> <li>• If a match occurs but the firewall filter does not specify an action, the packet is accepted.</li> <li>• 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.</li> <li>• 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.</li> <li>• If a match does not occur by the end of a firewall filter, the packet is discarded.</li> </ul>

## CHAPTER 2

# Introduction to Routing Policy

- [Understanding Routing Policies on page 15](#)
- [Importing and Exporting Routes on page 17](#)
- [Default Routing Policies and Actions on page 18](#)
- [Creating Routing Policies on page 20](#)
- [Configuring a Routing Policy on page 20](#)
- [Routing Policy Named Match Conditions on page 21](#)
- [Route-Based Match Conditions on page 21](#)
- [Protocol-Based Match Conditions on page 22](#)
- [Routing Policy Actions on page 23](#)
- [Routing Policy Terms on page 23](#)
- [Applying a Routing Policy on page 24](#)
- [Routing Protocol Support for Import and Export Policies on page 24](#)
- [Applying a Routing Policy to Routing Protocols on page 25](#)
- [Applying Export Policies to the Forwarding Table on page 25](#)
- [Evaluating a Routing Policy on page 26](#)
- [How a Routing Policy Is Evaluated on page 26](#)
- [How a Routing Policy Chain Is Evaluated on page 27](#)
- [How a Routing Policy Expression Is Evaluated on page 28](#)
- [How a Routing Policy Subroutine Is Evaluated on page 28](#)
- [Routing Policy Tests on page 30](#)

## Understanding Routing Policies

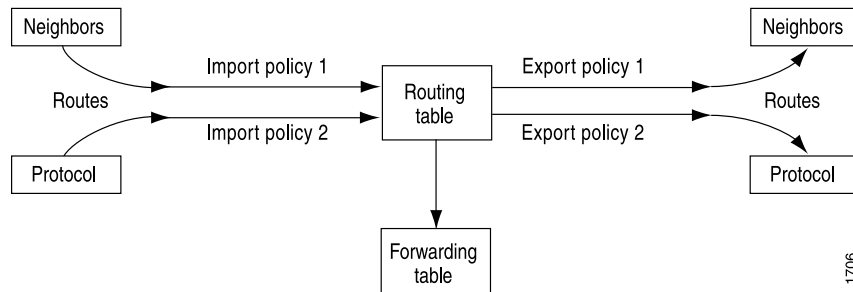
---

For some routing platform vendors, the flow of routes occurs between various protocols. If, for example, you want to configure redistribution from RIP to OSPF, the RIP process tells the OSPF process that it has routes that might be included for redistribution. In Junos OS, there is not much direct interaction between the routing protocols. Instead, there are central gathering points where all protocols install their routing information. These are the main unicast routing tables `inet.0` and `inet6.0`.

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.

As shown in [Figure 5 on page 16](#), you use import routing policies to control which routes are placed in the routing table, and export routing policies to control which routes are advertised from the routing table to neighbors.

**Figure 5: Importing and Exporting Routing Policies**



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 Junos OS 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.
- Manipulate the route characteristics as a routing protocol places the route in the routing table or advertises the route 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 is 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.

In Junos OS Release 9.5 and later, you can configure routing policies and certain routing policy objects in a dynamic database that is not subject to the same verification required by the standard configuration database. As a result, you can quickly commit these routing policies and policy objects, which can be referenced and applied in the standard configuration as needed. BGP is the only protocol to which you can apply routing policies that reference policies configured in the dynamic database. After a routing policy based on the dynamic database is configured and committed in the standard configuration, you can quickly make changes to existing routing policies by modifying policy objects in the dynamic database. Because Junos OS does not validate configuration changes to the dynamic database, when you use this feature, you should test and verify all configuration changes before committing them.

#### Related Documentation

- [Default Routing Policies and Actions on page 18](#)
- [Routing Tables Affected by Routing Policies on page 256](#)
- [Configuring Dynamic Routing Policies on page 76](#)
- [Example: Redistributing OSPF Routes into IS-IS](#)

## Importing and Exporting Routes


Two terms—*import* and *export*—explain how routes move between the routing protocols and the routing table (see [Figure 6 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 6: 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.



**NOTE:** By default, 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.

[“Protocols That Can Be Imported to and Exported from the Routing Table” on page 255](#) lists the routing protocols from which the routing table can import routes and to which the routing table can export routes. That topic 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 [“Default Import and Export Policies for Protocols” on page 257](#). 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 [“Protocols That Can Be Imported to and Exported from the Routing Table” on page 255](#).

---

## 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. [“Default Import and Export Policies for Protocols” on page 257](#) summarizes the default routing policies for each routing protocol that imports and exports routes. The actions in the default routing policies 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.

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 OS 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 [“Protocols That Can Be Imported to and Exported from the Routing Table” on page 255](#). 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 [“Protocol Support for Import and Export Policies” on page 261](#).

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.

The default import policy is always the same: accept all routes learned from the protocol. [“Default Import and Export Policies for Protocols” on page 257](#) includes information about the routing tables used by each protocol.

## Creating Routing Policies

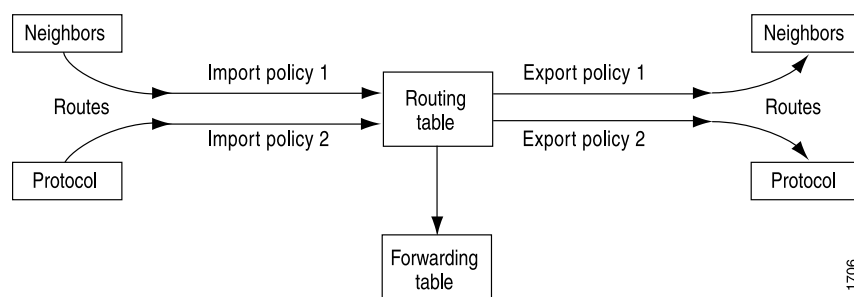
The following are typical circumstances under which you might want to preempt the default routing policies in the routing policy framework by creating your own routing policies:

- 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 7 on page 20](#), 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 7: 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 8 on page 21](#) shows the routing policy components, including the term.

**Figure 8: 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:

- [Routing Policy Match Conditions on page 258](#)
- [Routing Policy Named Match Conditions on page 21](#)
- [Routing Policy Actions on page 23](#)
- [Routing Policy Terms on page 23](#)
- [Applying a Routing Policy on page 24](#)

## Routing Policy 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 [“Routing Policy Match Conditions” on page 258](#).

## Route-Based Match Conditions

You can specify known invalid (“bad”) routes to ignore by specifying matches on destination prefixes. When specifying a destination prefix, you can specify an exact match with a specific route, or a less precise match by using match types. You can configure either a common reject action that applies to the entire list, or an action associated with each prefix.

Additionally, you can specify that “good” routes be processed in a particular way. For instance, you can group traffic from specific source or destination addresses into forwarding classes to be processed using the class of service (CoS) feature.

[Table 6 on page 22](#) lists route list match types.

Table 6: Route List Match Types

Match Type	Match Conditions
<b>exact</b>	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.
<b>longer</b>	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.
<b>orlonger</b>	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.
<b>prefix-length-range <i>prefix-length2-prefix-length3</i></b>	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.
<b>through <i>destination-prefix</i></b>	<p>All the following are true:</p> <ul style="list-style-type: none"> <li>• The route shares the same most-significant bits (described by <i>prefix-length</i>) of the first destination prefix.</li> <li>• 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.</li> <li>• The number of bits in the route's prefix length is less than or equal to the number of bits in the second prefix.</li> </ul> <p>You do not use the <b>through</b> match type in most routing policy configurations.</p>
<b>upto <i>prefix-length2</i></b>	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> .

- Related Documentation**
- [Routing Policy Match Conditions on page 258](#)
  - [Default Policies and Actions on page 8](#)
  - [Example: Rejecting Known Invalid Routes on page 86](#)
  - [Example: Grouping Source and Destination Prefixes into a Forwarding Class on page 106](#)
  - [Routing Policies Configuration Overview on page 10](#)

## Protocol-Based Match Conditions

You can specify a match condition for policies based on protocols by naming a protocol from which the route is learned or to which the route is being advertised. You can specify one of the following protocols: aggregate, BGP, direct, DVMRP, IS-IS, local, OSPF, PIM-dense, PIM-sparse, RIP, or static.

- Related Documentation**
- [Junos OS Feature Support Reference for SRX Series and J Series Devices](#)
  - [Routing Policy Match Conditions on page 258](#)

- [Default Policies and Actions on page 8](#)
- [Example: Injecting OSPF Routes into the BGP Routing Table](#)
- [Routing Policies Configuration Overview on page 10](#)

## Routing Policy 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 [“Configuring Actions in Routing Policy Terms” on page 58](#).

## Routing Policy Terms

---

A *term* is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term 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 double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of

**accept** or **reject** is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

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.

**Related  
Documentation**

- [Understanding Routing Policies on page 15](#)
- [Example: Creating a Routing Policy Term on page 85](#)

---

## Applying a Routing Policy

After defining a routing policy, as discussed in [“Routing Policy Match Conditions” on page 258](#) and [“Routing Policy Actions” on page 23](#), 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 Protocol Support for Import and Export Policies on page 24](#)
- [Protocol Support for Import and Export Policies on page 261](#)
- [Applying a Routing Policy to Routing Protocols on page 25](#)
- [Applying Export Policies to the Forwarding Table on page 25](#)

---

## Routing Protocol Support for Import and Export Policies

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. [“Protocol Support for Import and Export Policies” on page 261](#) summarizes the import and export policy support for each routing protocol. [“Protocols That Can Be Imported to and Exported from the Routing Table” on page 255](#) 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 RIPv6 only, you can apply import policies at the global and neighbor levels and export policies at a group level.

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 [“Protocols That Can Be Imported to and Exported from the Routing Table” on page 255](#). For information about the default routing policies for each routing protocol, see [“Default Import and Export Policies for Protocols” on page 257](#).

---

## Applying a Routing Policy 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.

---

## Applying Export Policies to the 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 Per-Packet Load Balancing” on page 186](#).

- Related Documentation**
- [Junos OS 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 26](#)
- [How a Routing Policy Chain Is Evaluated on page 27](#)
- [How a Routing Policy Expression Is Evaluated on page 28](#)
- [How a Routing Policy Subroutine Is Evaluated on page 28](#)

For specific information about how the various match conditions are evaluated, see [“Configuring Match Conditions in Routing Policy Terms” on page 51](#).

---

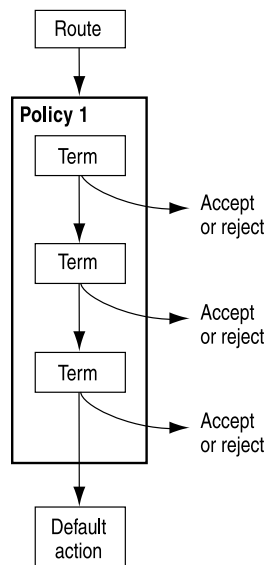
## How a Routing Policy Is Evaluated

[Figure 9 on page 27](#) 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 18](#).



Figure 9: Routing Policy Evaluation




## How a Routing Policy Chain Is Evaluated

Figure 10 on page 28 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 18](#).

Figure 10: 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 to Routes Exported from Routing Tables” on page 70](#).

## How a Routing Policy Subroutine Is Evaluated

Figure 11 on page 30 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 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 [“Configuring Subroutines in Routing Policy Match Conditions” on page 170](#). For more information about the default routing policies, see [“Default Routing Policies and Actions” on page 18](#).)

- 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 27](#).



**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 11: Routing Policy Subroutine Evaluation



## 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, BGP routes that match the prefix are advertised to 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.

## CHAPTER 3

# Introduction to Extended Match Conditions Configuration

- [Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions on page 31](#)
- [How BGP Communities and Extended Communities Are Evaluated in Routing Policy Match Conditions on page 32](#)

## Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions

---

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 enables you to perform actions on a group without having to elaborate upon each member. You can create a named community and include it in a routing policy with the **community** match condition, which is described in [“Configuring Match Conditions in Routing Policy Terms” on page 51](#). For a list of the actions that can be configured for communities, see [“Configuring Actions That Manipulate Route Characteristics” on page 60](#).

You can configure the standard community attribute and the extended communities attribute for inclusion in BGP update messages. The standard community attribute is four octets whereas the extended communities attribute is eight octets, providing a larger range 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 BGP community attribute format is **as-number:community-value**. The BGP extended communities attribute format has three fields: **type:administrator:assigned-number**.

When specifying community IDs for the standard community attribute, you can use UNIX-style regular expressions. Regular expressions are also supported for the extended communities attribute. The only exception is for VPN import policies (**vrf-import**), which do not support regular expressions for the extended communities attribute.



**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.

To use a BGP community or extended community as a routing policy match condition, you define the community and its members and then include the community in a match condition.

Junos OS supports the following standard:

- RFC 1997, *BGP Communities Attribute*

**Related  
Documentation**

- [Defining BGP Communities and Extended Communities for Use in Routing Policy Match Conditions on page 140](#)
- [Including BGP Communities and Extended Communities in Routing Policy Match Conditions on page 146](#)
- [How BGP Communities and Extended Communities Are Evaluated in Routing Policy Match Conditions on page 32](#)
- [Using Routing Policies to Prevent Advertisement of BGP Communities to Neighbors on page 147](#)
- [Examples: Configuring BGP Communities as Routing Policy Match Conditions on page 147](#)
- [Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS](#)

---

## How BGP Communities and Extended Communities Are Evaluated in Routing Policy Match Conditions

---

When you use BGP communities and extended communities as match conditions in a routing policy, the policy framework software evaluates them 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 identified by either a literal community value or a regular expression (for information about using regular expressions, see [“Using UNIX Regular Expressions in Community Names” on page 142](#)). 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 ];
}
```

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

For example:

```
community example1 members 100:100
community example2 members 100:1..
```

Given a route with a community value of 1100:100, this route matches **community example2** but not **example1**.

- 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.

## Multiple Matches

When multiple matches are found, label aggregation does not happen. Consider the following configuration:

```
family inet-vpn {
  unicast {
    aggregate-label {
      community community-name;
    }
  }
}

family inet-vpn {
  labeled-unicast {
    aggregate-label {
      community community-name;
    }
  }
}
```

Suppose, for instance, that two routes are received with community attributes **target:65000:1000** **origin:65200:2000** and that the community name is **"5....\*"**. In this case, both the extended community attributes, **target:65000:1000** and **origin:65200:2000** match the regular expression of the community name. In this case, label aggregation does not occur. In the following example, the **Label operation** field shows that the labels are not aggregated.

```
user@host> show route table VPN detail | match "^10 | Communities | Push"
10.1.1.0/30 (1 entry, 1 announced)
    Label operation: Push 101040
    Push 101040
    Communities: target:65000:1000 origin:65200:2000
10.1.1.4/30 (1 entry, 1 announced)
    Label operation: Push 101056
    Push 101056
    Communities: target:65000:1000 origin:65200:2000
```

You can resolve this issue in either of the following ways:

- Be more specific in the regular expression if the site-of-origin extended community attribute does not overlap with the target one.
- Specify the site of origin in the community name.

Both methods are shown in the following examples.

#### Be More Specific in the Regular Expression

```
user@host# set policy-options community community-name members "52....*"
user@host# commit
```

```
user@host> show route table VPN detail | match "^10 | Communities | Push"
10.1.1.0/30 (1 entry, 1 announced)
    Label operation: Push 101040
    Push 101040
    Communities: target:65000:1000 origin:65200:2000
10.1.1.4/30 (1 entry, 1 announced)
    Label operation: Push 101040
    Push 101040
    Communities: target:65000:1000 origin:65200:2000
```

#### Specify the Site of Origin in the Community Name

```
user@host# set policy-options community community-name members "origin:65....*"
user@host# commit
```

```
user@host> show route table VPN detail | match "^10 | Communities | Push"
10.1.1.0/30 (1 entry, 1 announced)
    Label operation: Push 101040
    Push 101040
    Communities: target:65000:1000 origin:65200:2000
10.1.1.4/30 (1 entry, 1 announced)
    Label operation: Push 101040
    Push 101040
    Communities: target:65000:1000 origin:65200:2000
```

## Extended Community Type

The extended community type is not taken into account by regular expressions. Consider, for instance, the following community attributes and community name.



Communities:

- 5200:1000
- **target:65000:1000**
- **origin:65200:2000**

Community attribute:

- community-name members "5....:"

In this case, both extended community attribute, **5200:1000** and the extended community attribute, **origin:65200:2000**, match the regular expression of the community name. Therefore, the label aggregation does not occur, as shown here:

```
user@host> show route table VPN detail | match "^10 | Communities | Push"
10.1.1.0/30 (1 entry, 1 announced)
    Label operation: Push 101040
    Push 101040
    Communities: 5200:1000 target:65000:1000 origin:65200:2000
10.1.1.4/30 (1 entry, 1 announced)
    Label operation: Push 101056
    Push 101056
    Communities: 5200:1000 target:65000:1000 origin:65200:2000
```

You can resolve this issue by using a more specific regular expression. For example, you can use the anchor character (^) to bind the location of the digits, as shown here:

```
user@host# set policy-options community community-name members "^5....:"
user@host# commit

user@host> show route table VPN detail | match "^10 | Communities | Push"
10.1.1.0/30 (1 entry, 1 announced)
    Label operation: Push 101040
    Push 101040
    Communities: 5200:1000 target:65000:1000 origin:65200:2000
10.1.1.4/30 (1 entry, 1 announced)
    Label operation: Push 101040
    Push 101040
    Communities: 5200:1000 target:65000:1000 origin:65200:2000
```

## Multiple Communities Are Matched with Ex-OR Logic

This differs from the AND matching logic used for non-extended communities in BGP.

If, for instance, four routes are received with two sets of community attributes, the regular expression might match both community attributes. Consider the following example:

- Communities—5200:1000 target:65000:1000
- Communities—target:65000:1000 origin:65200:2000
- Community attribute—community community-name member [ "^5....:" origin:65.\*.\* ]

Both labels are aggregated, as shown here:

```
user@host> show route table VPN detail | match "^10 | Communities | Push"
```

```
10.1.1.0/30 (1 entry, 1 announced)
    Label operation: Push 101040
    Push 101040
    Communities: target:65000:1000 origin:65200:2000
10.1.1.4/30 (1 entry, 1 announced)
    Label operation: Push 101040
    Push 101040
    Communities: target:65000:1000 origin:65200:2000

10.1.1.16/30 (1 entry, 1 announced)
    Label operation: Push 121104
    Push 101104
    Communities: 5200:1000 target:65000:1000
10.1.1.20/30 (1 entry, 1 announced)
    Label operation: Push 121104
    Push 101104
    Communities: 5200:1000 target:65000:1000
```

A more complete example of community values is shown here:

```
user@host> show policy-options community community-name
members [ "(^1...:*)" | (^3...:*)" | (^4...:*)" origin:2.*:* origin:3.*:* origin:6.*:*
]
```

This regular expression matches community values starting with 1, 3, or 4, and matches extended community values of type origin whose administrative value starts with 2, 3, or 6.

## CHAPTER 4

# Introduction to Extended Actions Configuration

- [Overview of Per-Packet Load Balancing on page 37](#)

## Overview of Per-Packet Load Balancing

---

By default, when there are multiple equal-cost paths to the same destination for the active route, the Junos OS 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 OS 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 the 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.
- On some platforms, you can increase the number of paths by using the **chassis maximum-ecmp** statement. With this statement, you can change the maximum number of equal-cost load-balanced paths to 32 or 64.

### Related Documentation

- [Configuring Per-Packet Load Balancing on page 186](#)
- [Configuring Load Balancing Based on MPLS Labels on DPC I-Chip-Based Hardware on page 189](#)

- [Configuring Load Balancing for Ethernet Pseudowires on page 193](#)
- [Configuring Load Balancing Based on MAC Addresses on page 194](#)
- [Configuring VPLS Load Balancing Based on IP and MPLS Information on page 195](#)
- [Configuring VPLS Load Balancing on MX Series 3D Universal Edge Routers on page 196](#)

## CHAPTER 5

# Introduction to Traffic Sampling, Forwarding, and Monitoring Overview

- [Traffic Sampling, Forwarding, and Monitoring Overview on page 39](#)

## Traffic Sampling, Forwarding, and Monitoring 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. In Junos OS Release 9.0 and later, 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 routers 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 6

# Introduction to Traffic Forwarding and Monitoring Configuration

- [Per-Flow and Per-Prefix Load Balancing Overview on page 41](#)

### Per-Flow and Per-Prefix Load Balancing Overview

By default, when there are multiple equal-cost paths to the same destination, the Junos OS chooses one of the next-hop addresses at random.

On all M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers, you have the additional option of configuring per-prefix load balancing based on a specified hash value that enables the router to elect a next hop independently of the route chosen by other routers.

On the M120, M320, and MX Series routers only, you have the additional option of enabling per-flow load balancing based on a unique, load-balance hash value for each Packet Forwarding Engine slot.

#### **Related Documentation**

- [Configuring Per-Prefix Load Balancing on page 238](#)
- [Configuring Per-Flow Load Balancing Based on Hash Values on page 239](#)





## PART 2

# Configuration

- [Routing Policy Configuration Statements on page 45](#)
- [Routing Policy Configuration on page 49](#)
- [Extended Match Conditions Configuration on page 133](#)
- [Extended Actions Configuration on page 177](#)
- [Traffic Sampling Configuration on page 211](#)
- [Traffic Forwarding and Monitoring Configuration on page 225](#)



## CHAPTER 7

# Routing Policy Configuration Statements

- [Configuring a Routing Policy on page 45](#)
- [Minimum Routing Policy Configuration on page 46](#)
- [Minimum Routing Policy Chain Configuration on page 46](#)
- [Minimum Subroutine Configuration on page 47](#)

### Configuring a Routing Policy

---

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

```
policy-options {  
  as-path 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 source-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 {
    export [ policy-names ];
    import [ policy-names ];
  }
}
```

---

## 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 source-prefix match-type <actions>;
      }
      to {
        match-conditions;
      }
      then actions;
    }
  }
}
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 source-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 source-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 *protocol-name* (export | import)]** hierarchy levels, you can specify 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 source-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 8

# Routing Policy Configuration

- [Defining Routing Policies on page 50](#)
- [Configuring Match Conditions in Routing Policy Terms on page 51](#)
- [Configuring Actions in Routing Policy Terms on page 58](#)
- [Applying Routing Policies and Policy Chains to Routing Protocols on page 70](#)
- [Applying Policy Expressions to Routes Exported from Routing Tables on page 70](#)
- [Applying Routing Policies to the Forwarding Table on page 75](#)
- [Configuring Dynamic Routing Policies on page 76](#)
- [Forwarding Packets to the Discard Interface on page 81](#)
- [Testing Routing Policies on page 82](#)
- [Example: Creating a Routing Policy on page 83](#)
- [Routing Policy Examples on page 84](#)
- [Example: Creating a Routing Policy Term on page 85](#)
- [Example: Rejecting Known Invalid Routes on page 86](#)
- [Examples: Configuring Static Routes on page 87](#)
- [Example: Defining a Routing Policy from BGP to IS-IS on page 97](#)
- [Example: Using Routing Policy to Set a Preference on page 98](#)
- [Example: Importing and Exporting Access and Access-Internal Routes in a Routing Policy on page 98](#)
- [Example: Exporting Routes to IS-IS on page 99](#)
- [Example: Applying Export and Import Policies to BGP Peer Groups on page 99](#)
- [Example: Defining a Routing Policy Based on the Number of BGP Communities on page 100](#)
- [Example: Applying a Prefix to Routes Learned from a Peer on page 100](#)
- [Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS on page 100](#)
- [Example: Redistributing OSPF Routes into BGP on page 101](#)
- [Example: Exporting Direct Routes Into IS-IS on page 102](#)
- [Example: Exporting Internal IS-IS Level 1 Routes to Level 2 on page 102](#)
- [Example: Exporting IS-IS Level 2 Routes to Level 1 on page 102](#)

- [Example: Assigning Different Forwarding Next-Hop LSPs to Different Destination Prefixes on page 103](#)
- [Example: Grouping Destination Prefixes on page 104](#)
- [Example: Grouping Source Prefixes on page 105](#)
- [Example: Grouping Source and Destination Prefixes into a Forwarding Class on page 106](#)
- [Example: Accepting Routes with Specific Destination Prefixes on page 108](#)
- [Example: Accepting Routes from BGP with a Specific Destination Prefix on page 108](#)
- [Example: Using Routing Policy in an ISP Network on page 109](#)
- [Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Switches on page 121](#)
- [Understanding ECMP Flow-Based Forwarding on page 126](#)
- [Example: Configuring ECMP Flow-Based Forwarding on page 127](#)

---

## 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 source-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]**

***policy-name*** specifies the policy name and must be unique in the configuration. 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 (" ").

***term-name*** specifies the name of a term in the policy and must be unique in that policy. It 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 (" ").



A policy statement can include multiple terms, including a unnamed term which must be the final term in the policy. To configure an unnamed term, omit the **term** statement when defining match conditions and actions. However, we recommend that you name all terms.

To list the routing policies by **policy-statement** *policy-name* in alphabetical order, enter the **show policy-options** configuration command.

For information about configuring the components of a term, see the following sections:

- [Configuring Match Conditions in Routing Policy Terms on page 51](#)
- [Configuring Actions in Routing Policy Terms on page 58](#)

## Configuring Match Conditions in Routing Policy Terms

Each term in a routing policy can include two statements, **from** and **to**, to define the conditions that a route must match for the policy to apply:

```
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;
}
```

You can include these statements at the following hierarchy levels:

- [edit policy-options **policy-statement** *policy-name* term *term-name*]
- [edit logical-systems *logical-system-name* policy-options **policy-statement** *policy-name* term *term-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 [“Applying Routing Policies and Policy Chains to Routing Protocols” on page 70](#).

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 defined in [Table 7 on page 52](#) are effectively a logical AND operation. Matching in prefix lists and route lists is handled differently. They are effectively a logical OR operation. For more information about how matching occurs for prefix lists and route lists, including how they are evaluated, see “[Configuring Prefix Lists for Use in Routing Policy Match Conditions](#)” on page 151 and “[Configuring Route Lists for Use in Routing Policy Match Conditions](#)” on page 155. If you configure a policy that includes some combination of route filters, prefix lists, and source address filters, they are evaluated according to a logical OR operation or a longest-route match lookup.

[Table 7 on page 52](#) 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 7 on page 52](#) 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. The standard match conditions include criteria that are defined within a routing policy and are less complex than the extended match conditions.

**Table 7: Routing Policy Match Conditions**

Match Condition	Match Condition Category	from Statement Description	to Statement Description
<b>aggregate-contributor</b>	Standard	Match routes that are contributing to a configured aggregate. This match condition can be used to suppress a contributor in an aggregate route.	
<b>area <i>area-id</i></b>	Standard	(Open Shortest Path First [OSPF] only) Area identifier.  In a <b>from</b> statement used with an <b>export</b> policy, match a route learned from the specified OSPF area when exporting OSPF routes into other protocols.	

Table 7: Routing Policy Match Conditions (*continued*)

Match Condition	Match Condition Category	from Statement Description	to Statement Description
<b>as-path</b> <i>name</i>	Extended	(Border Gateway Protocol [BGP] only) Name of an AS path regular expression. For more information, see “ <a href="#">Configuring AS Path Regular Expressions to Use as Routing Policy Match Conditions</a> ” on page 133.	
<b>as-path-group</b> <i>group-name</i>	Extended	(BGP only) Name of an AS path group regular expression. For more information, see “ <a href="#">Configuring AS Path Regular Expressions to Use as Routing Policy Match Conditions</a> ” on page 133.	
<b>color</b> <i>preference</i> <b>color2</b> <i>preference</i>	Standard	Color value. You can specify preference values ( <b>color</b> and <b>color2</b> ) that are finer-grained than those specified in the <b>preference</b> and <b>preference2</b> 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.	
<b>community-count</b> <i>value (equal   orhigher   orlower)</i>	Standard	<p>(BGP only) Number of community entries required for a route to match. The count value can be a number in the range of 0 through 1,024. Specify one of the following options:</p> <ul style="list-style-type: none"> <li><b>equal</b>—The number of communities must equal this value to be considered a match.</li> <li><b>orhigher</b> —The number of communities must be greater than or equal to this value to be considered a match.</li> <li><b>orlower</b>—The number of communities must be less than or equal to this value to be considered a match.</li> </ul> <p><b>NOTE:</b> If you configure multiple <b>community-count</b> statements, the matching is effectively a logical AND operation.</p> <p><b>NOTE:</b> The <b>community-count</b> attribute only works with standard communities. It does not work with extended communities.</p>	
<b>community</b> [ <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 “ <a href="#">Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions</a> ” on page 31.	
<b>external</b> [ <i>type</i> <i>metric-type</i> ]	Standard	<p>(OSPF and IS-IS only) Match IGP external routes. For IS-IS routes, the <b>external</b> condition also matches routes that are exported from one IS-IS level to another. The <b>type</b> keyword is optional and is applicable only to OSPF external routes. When you do not specify <b>type</b>, the <b>external</b> condition matches all IGP external (OSPF and IS-IS) routes. When you specify <b>type</b>, the <b>external</b> condition matches only OSPF external routes with the specified OSPF metric type. The metric type can either be 1 or 2.</p> <p>To match BGP external routes, use the <b>route-type</b> match condition.</p>	

Table 7: Routing Policy Match Conditions (*continued*)

Match Condition	Match Condition Category	from Statement Description	to Statement Description
<b>family</b> <i>family-name</i>	Standard	<p>Name of an address family. Match the address family of the route. Depending on your device and configuration, <b>family-name</b> can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>inet</b>—IP version 4 (IPv4) traffic</li> <li>• <b>inet-mdt</b>—IPv4 multicast distribution tree (MDT) traffic</li> <li>• <b>inet-mvpn</b>—IPv4 multicast virtual private network (MVPN) traffic</li> <li>• <b>inet-vpn</b>—IPv4 VPN traffic</li> <li>• <b>inet6</b>—IP version 6 (IPv6) traffic</li> <li>• <b>inet6-mvpn</b>—IPv6 MVPN traffic</li> <li>• <b>inet6-vpn</b>—IPv6 VPN traffic</li> <li>• <b>iso</b>—IS-IS traffic</li> <li>• <b>route-target</b>—BGP route target filtering routes for VPN traffic</li> </ul> <p>Default setting is <b>inet</b>.</p>	
<b>instance</b> <i>instance-name</i>	Standard	<p>Name of one or more routing instances.</p> <p>Match a route learned from one of the specified instances.</p>	<p>Name of one or more routing instances.</p> <p>Match a route to be advertised over one of the specified instances.</p>
<b>interface</b> <i>interface-name</i>	Standard	<p>Name or IP address of one or more routing device interfaces. Do not use this qualifier with protocols that are not interface-specific, such as IBGP.</p> <p>Match a route learned from one of the specified interfaces. Direct routes match routes configured on the specified interface.</p>	<p>Name or IP address of one or more routing device interfaces. Do not use this qualifier with protocols that are not interface-specific, such as IBGP.</p> <p>Match a route to be advertised from one of the specified interfaces.</p>
<b>level</b> <i>level</i>	Standard	<p>(Intermediate System-to-Intermediate System [IS-IS] only) IS-IS level.</p> <p>Match a route learned from a specified level.</p>	<p>(IS-IS only) IS-IS level.</p> <p>Match a route to be advertised to a specified level.</p>
<b>local-preference</b> <i>value</i>	Standard	(BGP only) BGP local preference ( <b>LOCAL_PREF</b> ) attribute. The preference value can be a number in the range 0 through 4,294,967,295 ( $2^{32} - 1$ ).	
<b>metric</b> <i>metric</i> <b>metric2</b> <i>metric3</i> <i>metric</i> <i>metric4</i> <i>metric</i>	Standard	<p>Metric value. You can specify up to four metric values, starting with <b>metric</b> (for the first metric value) and continuing with <b>metric2</b>, <b>metric3</b>, and <b>metric4</b>.</p> <p>(BGP only) <b>metric</b> corresponds to the multiple exit discriminator (MED), and <b>metric2</b> corresponds to the interior gateway protocol (IGP) metric if the BGP next hop runs back through another route.</p>	

Table 7: Routing Policy Match Conditions (*continued*)

Match Condition	Match Condition Category	from Statement Description	to Statement Description
<b>multicast-scoping</b> ( <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 <b>orhigher</b> to match on a scope and numerically higher scopes, or <b>orlower</b> to match on a scope and numerically lower scopes. For more information, see the Multicast Protocols Configuration Guide.</p> <p>You can apply this scoping policy to the routing table by including the <b>scope-policy</b> statement at the <b>[edit routing-options]</b> hierarchy level.</p> <p>The <b>number</b> 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"> <li>• <b>node-local</b> (value=1)—No corresponding prefix</li> <li>• <b>link-local</b> (value=2)—Corresponding prefix <b>224.0.0.0/24</b></li> <li>• <b>site-local</b> (value=5)—No corresponding prefix</li> <li>• <b>global</b> (value=14)—Corresponding prefix <b>224.0.1.0</b> through <b>238.255.255.255</b></li> <li>• <b>organization-local</b> (value=8)—Corresponding prefix <b>239.192.0.0/14</b></li> </ul>	
<b>neighbor address</b>	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><b>NOTE:</b> The <b>neighbor address</b> 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 <b>to neighbor</b> produces the same result as specifying <b>from neighbor</b>.</p> <p>For BGP export policies, specifying the <b>neighbor</b> match condition has no effect and is ignored.</p> <p>For all other protocols, the <b>to</b> statement matches the neighbor to which the advertisement is sent.</p> <p><b>NOTE:</b> The <b>neighbor address</b> match condition is not valid for the Routing Information Protocol (RIP).</p>
<b>next-hop</b> [ <i>addresses</i> ]	Standard	One or more next-hop addresses specified in the routing information for a particular route. A next-hop address cannot include a netmask. For BGP routes, matches are performed against each protocol next hop.	
<b>next-hop-type merged</b>	Standard	LDP generates a next hop based on RSVP and IP next hops available to use, combined with forwarding-class mapping.	You cannot specify this match condition.
<b>nlri-route-type</b>	Standard	<p>Route type from NLRI 1 through NLRI 10.</p> <p>Multiple route types can be specified in a single policy.</p>	

Table 7: Routing Policy Match Conditions (*continued*)

Match Condition	Match Condition Category	from Statement Description	to Statement Description
<b>origin value</b>	Standard	(BGP only) BGP origin attribute, which is the origin of the AS path information. The value can be one of the following: <ul style="list-style-type: none"> <li>• <b>egp</b>—Path information originated in another AS.</li> <li>• <b>igp</b>—Path information originated within the local AS.</li> <li>• <b>incomplete</b>—Path information was learned by some other means.</li> </ul>	
<b>policy [ <i>policy-name</i> ]</b>	Extended	Name of a policy to evaluate as a subroutine.  For information about this extended match condition, see <a href="#">“Configuring Subroutines in Routing Policy Match Conditions” on page 170.</a>	
<b>preference preference preference2 preference</b>	Standard	Preference value. You can specify a primary preference value ( <b>preference</b> ) and a secondary preference value ( <b>preference2</b> ). 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 <b>color</b> and <b>color2</b> match conditions in this table.	
<b>prefix-list prefix-list-name ip-addresses</b>	Extended	Named list of IP addresses. You can specify an exact match with incoming routes.  For information about this extended match condition, see <a href="#">“Configuring Prefix Lists for Use in Routing Policy Match Conditions” on page 151.</a>	You cannot specify this match condition.
<b>prefix-list-filter prefix-list-name match-type</b>	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 <a href="#">“Configuring Prefix Lists for Use in Routing Policy Match Conditions” on page 151.</a>	You cannot specify this match condition.
<b>protocol protocol</b>	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: <b>access</b> , <b>access-internal</b> , <b>aggregate</b> , <b>bgp</b> , <b>direct</b> , <b>dvmrp</b> , <b>isis</b> , <b>l2vpn</b> , <b>ldp</b> , <b>local</b> , <b>msdp</b> , <b>ospf</b> , <b>ospf2</b> , <b>ospf3</b> , <b>pim</b> , <b>rip</b> , <b>ripng</b> , <b>route-target</b> , <b>rsvp</b> , or <b>static</b> .  <b>NOTE:</b> The <b>ospf2</b> statement matches on OSPFv2 routes. The <b>ospf3</b> statement matches on OSPFv3 routes. The <b>ospf</b> statement matches on both OSPFv2 and OSPFv3 routes.  For more information about access routes and access-internal routes, see <a href="#">“Example: Importing and Exporting Access and Access-Internal Routes in a Routing Policy” on page 98.</a>	

Table 7: Routing Policy Match Conditions (*continued*)

Match Condition	Match Condition Category	from Statement Description	to Statement Description
<b>rib <i>routing-table</i></b>	Standard	<p>Name of a routing table. The value of <b><i>routing-table</i></b> can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>inet.0</b>—Unicast IPv4 routes</li> <li>• <b><i>instance-name</i> inet.0</b>—Unicast IPv4 routes for a particular routing instance</li> <li>• <b>inet.1</b>—Multicast IPv4 routes</li> <li>• <b>inet.2</b>—Unicast IPv4 routes for multicast reverse-path forwarding (RPF) lookup</li> <li>• <b>inet.3</b>—MPLS routes</li> <li>• <b>mpls.0</b>—MPLS routes for label-switched path (LSP) next hops</li> <li>• <b>inet6.0</b>—Unicast IPv6 routes</li> </ul>	
<b>route-filter</b> <b><i>destination-prefix</i></b> <b><i>match-type</i></b> <b>&lt;<i>actions</i>&gt;</b>	Extended	<p>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 <a href="#">“Configuring Route Lists for Use in Routing Policy Match Conditions” on page 155</a>.</p>	You cannot specify this match condition.
<b>route-type <i>value</i></b>	Standard	<p>Type of BGP route. The value can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>external</b>—External route.</li> <li>• <b>internal</b>—Internal route.</li> </ul> <p>To match IGP external routes, use the <b>external</b> match condition.</p>	
<b>rtf-prefix-list <i>name</i></b> <b><i>route-targets</i></b>	Extended	<p>(BGP only) Named list of route target prefixes for BGP route target filtering and proxy BGP route target filtering.</p> <p>For information about this extended match condition, see Example: Configuring Proxy BGP Route Target Filtering.</p>	You cannot specify this match condition.
<b>source-address-filter</b> <b><i>destination-prefix</i></b> <b><i>match-type</i></b> <b>&lt;<i>actions</i>&gt;</b>	Extended	<p>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 <a href="#">“Configuring Route Lists for Use in Routing Policy Match Conditions” on page 155</a>.</p>	You cannot specify this match condition.

Table 7: Routing Policy Match Conditions (*continued*)

Match Condition	Match Condition Category	from Statement Description	to Statement Description
<b>state (active   inactive)</b>	Standard	(BGP export only) Match on the following types of advertised routes: <ul style="list-style-type: none"> <li>• <b>active</b>—An active BGP route</li> <li>• <b>inactive</b>—A route advertised to internal BGP peers as the best external path even if the best path is an internal route</li> <li>• <b>inactive</b>—A route advertised by BGP as the best route even if the routing table did not select it to be an active route</li> </ul>	
<b>tag string tag2 string</b>	Standard	<p>Tag value. You can specify two tag strings: <b>tag</b> (for the first string) and <b>tag2</b>. These values are local to the router and can be set on configured routes or by using an import routing policy.</p> <p>You can specify multiple tags under one match condition by including the tags within a bracketed list. For example: <b>from tag [ tag1 tag2 tag3 ]</b>;</p> <p>For OSPF routes, the <b>tag</b> action sets the 32-bit tag field in OSPF external link-state advertisement (LSA) packets.</p> <p>For IS-IS routes, the <b>tag</b> action sets the 32-bit flag in the IS-IS IP prefix type length values. (TLV).</p> <p>OSPF stores the INTERNAL route's OSPF area ID in the <b>tag2</b> attribute. However, for EXTERNAL routes, OSPF does not store anything in the <b>tag2</b> attribute.</p> <p>You can configure a policy term to set the <b>tag2</b> value for a route. If the route, already has a <b>tag2</b> value (for example, an OSPF route that stores area id in tag2), then the original <b>tag2</b> value is overwritten by the new value.</p> <p>When the policy contains the "from area" match condition, for internal OSPF routes, where <b>tag2</b> is set, based on the OSPF area- ID, the evaluation is conducted to compare the <b>tag2</b> attribute with the area ID. For external OSPF routes that do not have the <b>tag2</b> attribute set, the match condition fails.</p>	
<b>validation-database</b>	Standard	<p>When BGP origin validation is configured, triggers a lookup in the route validation database to determine if the route prefix is valid, invalid, or unknown. The route validation database contains route origin authorization (ROA) records that map route prefixes to expected originating autonomous systems (ASs). This prevents the accidental advertisement of invalid routes.</p> <p>See Example: Configuring Origin Validation for BGP.</p>	

**Related Documentation** • [Routing Policy Examples on page 84](#)

## Configuring Actions in Routing Policy Terms

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

```
then {
    actions;
```



```
}
```

You can include this statement at the following hierarchy levels:

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

If a term does not have **from** and **to** statements, all routes are considered to match, and the actions apply to them all. For information about the **from** and **to** statements, see [“Configuring Match Conditions in Routing Policy Terms” on page 51](#).

You can specify one or more actions in the **then** statement. 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.

The **then** statement is optional. If you omit it, 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 18](#).

The following sections discuss the following actions:

- [Configuring Flow Control Actions on page 60](#)
- [Configuring Actions That Manipulate Route Characteristics on page 60](#)
- [Configuring the Default Action in Routing Policies on page 67](#)

- [Configuring a Final Action in Routing Policies on page 68](#)
- [Logging Matches to a Routing Policy Term on page 69](#)
- [Configuring Separate Actions for Routes in Route Lists on page 69](#)

## Configuring Flow Control Actions

[Table 8 on page 60](#) lists the flow control actions. You can specify one of these actions along with the trace action or one or more of the actions that manipulate route characteristics (see [“Configuring Actions That Manipulate Route Characteristics” on page 60](#)).

**Table 8: Flow Control Actions**

Flow Control Action	Description
<b>accept</b>	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.
<b>default-action accept</b>	Accept and override any action intrinsic to the protocol. This is a nonterminating policy action.
<b>reject</b>	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.
<b>default-action reject</b>	Reject and override any action intrinsic to the protocol. This is a nonterminating policy action.
<b>next term</b>	<p>Skip to and evaluate the next term in the same routing policy. Any accept or reject action specified in the <b>then</b> statement is skipped. Any actions in the <b>then</b> statement that manipulate route characteristics are applied to the route.</p> <p><b>next term</b> is the default control action if a match occurs and you do not specify a flow control action.</p>
<b>next policy</b>	<p>Skip to and evaluate the next routing policy. Any accept or reject action specified in the <b>then</b> statement is skipped. Any actions in the <b>then</b> statement that manipulate route characteristics are applied to the route.</p> <p><b>next policy</b> 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>

## Configuring Actions That Manipulate Route Characteristics

You can specify one or more of the actions listed in [Table 9 on page 61](#) to manipulate route characteristics.

Table 9: Actions That Manipulate Route Characteristics

Action	Description
<b>as-path-prepend <i>as-path</i></b>	<p>(BGP only) Affix one or more AS numbers at the beginning of the AS path. If specifying more than one AS number, enclose 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 <a href="#">“Prepending AS Numbers to BGP AS Paths” on page 177</a>.</p> <p>In Junos OS Release 9.1 and later, you can specify 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>, as well as the 2-byte AS numbers that are supported in earlier releases of the Junos OS.</p>
<b>as-path-expand last-as count <i>n</i></b>	<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>
<b>class <i>class-name</i></b>	<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 Junos OS Class of Service Configuration Guide.</p>
<b>color <i>preference</i> color2 <i>preference</i></b>	<p>Set the preference value to the specified value. The <b>color</b> and <b>color2</b> preference values are even more fine-grained than those specified in the <b>preference</b> and <b>preference2</b> actions. The color value can be a number in the range from 0 through 4,294,967,295 (<math>2^{32} - 1</math>). A lower number indicates a more preferred route.</p> <p>If you set the preference with the <b>color</b> action, the value is internal to the Junos OS and is not transitive.</p>
<b>color (add   subtract) <i>number</i></b> <b>color2 (add   subtract) <i>number</i></b>	<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 (<math>2^{32} - 1</math>), the value is set to <math>2^{32} - 1</math>. 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>
<b>community (+   add) [ <i>names</i> ]</b>	<p>(BGP only) Add the specified communities to the set of communities in the route. For more information, see <a href="#">“Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions” on page 31</a>.</p>
<b>community (–   delete) [ <i>names</i> ]</b>	<p>(BGP only) Delete the specified communities from the set of communities in the route. For more information, see <a href="#">“Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions” on page 31</a>.</p>
<b>community (=   set) [ <i>names</i> ]</b>	<p>(BGP only) Replace any communities that were in the route in with the specified communities. For more information, see <a href="#">“Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions” on page 31</a>.</p>
<b>cos-next-hop-map <i>map-name</i></b>	<p>Set CoS-based next-hop map in forwarding table.</p>

Table 9: Actions That Manipulate Route Characteristics (*continued*)

Action	Description
<b>damping</b> <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 Junos OS Routing Protocols Configuration Guide, and you must create a named list of parameters as described in <a href="#">“Using Routing Policies to Damp BGP Route Flapping” on page 181</a>.</p>
<b>destination-class</b> <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"> <li>• Configure group destination prefixes by configuring a routing policy; see <a href="#">“Defining Routing Policies” on page 50</a> and <a href="#">“Routing Policy Examples” on page 84</a>.</li> <li>• Apply that routing policy to the forwarding table with the corresponding destination class; see <a href="#">“Applying Routing Policies to the Forwarding Table” on page 75</a>.</li> <li>• Enable packet counting on one or more interfaces by including the <b>destination-class-usage</b> statement at the <b>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting]</b> hierarchy level (see the Junos OS Class of Service Configuration Guide). See <a href="#">“Routing Policy Examples” on page 84</a>.</li> <li>• View the output by using one of the following commands: <b>show interfaces destination-class (all   <i>destination-class-name</i> <i>logical-interface-name</i>)</b>, <b>show interfaces <i>interface-name</i> extensive</b>, or <b>show interfaces <i>interface-name</i> statistics</b> (see the Junos OS Operational Mode Commands).</li> <li>• To configure a packet count based on the source address, use the <b>source-class</b> statement described in this table.</li> </ul>
<b>external type</b> <i>metric</i>	Set the external metric type for routes exported by OSPF. You must specify the keyword <b>type</b> .
<b>forwarding-class</b> <i>forwarding-class-name</i>	<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"> <li>• Configure group prefixes by configuring a routing policy; see <a href="#">“Defining Routing Policies” on page 50</a> and <a href="#">“Routing Policy Examples” on page 84</a>.</li> <li>• Apply that routing policy to the forwarding table with the corresponding forwarding class; see <a href="#">“Applying Routing Policies to the Forwarding Table” on page 75</a>.</li> <li>• Enable packet counting on one or more interfaces by using the procedure described in either the <b>destination-class</b> or <b>source-class</b> actions defined in this table.</li> </ul>
<b>install-nexthop</b> <strict> <i>lsp-name</i>	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 <b>strict</b> 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.
<b>install-to-fib</b>	For PTX Series switches only, override the default BGP routing policy. For more information, see <a href="#">“Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Switches” on page 121</a> .

Table 9: Actions That Manipulate Route Characteristics (*continued*)

Action	Description
<b>load-balance per-packet</b>	(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 <a href="#">“Configuring Per-Packet Load Balancing” on page 186</a> .
<b>load-balance per-prefix</b>	For PTX Series switches only, override the default per-packet load balancing routing policy for BGP. For more information, see <a href="#">“Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Switches” on page 121</a> .
<b>local-preference value</b>	(BGP only) Set the BGP local preference ( <b>LOCAL_PREF</b> ) attribute. The preference value can be a number in the range from 0 through 4,294,967,295 ( $2^{32} - 1$ ).
<b>local-preference (add   subtract) number</b>	<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 (<math>2^{32} - 1</math>), the value is set to <math>2^{32} - 1</math>. 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>
<b>map-to-interface (interface-name   self)</b>	<p>Sets the <b>map-to-interface</b> value which is similar to existing metric or tag actions. The <b>map-to-interface</b> action requires you to specify one of the following:</p> <ul style="list-style-type: none"> <li>A logical interface (for example, ge-0/0/0.0). The logical interface can be any interface that multicast currently supports, including VLAN and aggregated Ethernet interfaces.</li> </ul> <p><b>NOTE:</b> If you specify a physical interface as the <b>map-to-interface</b> (for example, ge-0/0/0), a value of .0 is appended to physical interface to create a logical interface.</p> <ul style="list-style-type: none"> <li>The keyword <b>self</b>. The <b>self</b> keyword specifies that multicast data packets are sent on the same interface as the control packets and no mapping occurs.</li> </ul> <p>If no term matches, then no multicast data packets are sent.</p>
<b>metric metric metric2 metric metric3 metric metric4 metric</b>	<p>Set the metric. You can specify up to four metric values, starting with <b>metric</b> (for the first metric value) and continuing with <b>metric2</b>, <b>metric3</b>, and <b>metric4</b>.</p> <p>(BGP only) <b>metric</b> corresponds to the MED, and <b>metric2</b> corresponds to the IGP metric if the BGP next hop loops through another router.</p>
<b>metric (add   subtract) number metric2 (add   subtract) number metric3 (add   subtract) number metric4 (add   subtract) number</b>	<p>Change the metric value by the specified amount. If an addition operation results in a value that is greater than 4,294,967,295 (<math>2^{32} - 1</math>), the value is set to <math>2^{32} - 1</math>. 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>

Table 9: Actions That Manipulate Route Characteristics (*continued*)

Action	Description
<b>metric expression (metric multiplier <i>x</i> offset <i>a</i>   metric2 multiplier <i>y</i> offset <i>b</i>)</b>	<p>Calculate a metric based on the current values of <b>metric</b> and <b>metric2</b>.</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 <b>metric</b> and <b>metric2</b> are the current input values. Metric multipliers are limited in range to eight significant digits.</p>
<b>metric (igp   minimum-igp) site-offset</b>	<p>(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.</p>

Table 9: Actions That Manipulate Route Characteristics (*continued*)

Action	Description
<b>next-hop</b> ( <i>address</i>   <b>discard</b>   <b>next-table</b> <i>table-name</i>   <b>peer-address</b>   <b>reject</b>   <b>self</b> )	<p>Set the next-hop address. 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 <b>self</b>, the next-hop address is replaced by one of the local routing device'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 routing device cannot install routes with itself as the next hop.</p> <p>If you specify <b>peer-address</b>, 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 routing device or another directly connected routing device.</p> <p>If you specify <b>discard</b>, the next-hop address is replaced by a discard next hop.</p> <p>If you specify <b>next-table</b>, the routing device performs a forwarding lookup in the specified table.</p> <p>If you use the <b>next-table</b> action, the configuration must include a term qualifier that specifies a different table than the one specified in the <b>next-table</b> action. In other words, the term qualifier in the <b>from</b> statement must exclude the table in the <b>next-table</b> action. In the following example, the first term contains <b>rib vrf-customer2.inet.0</b> as a matching condition. The action specifies a next-hop in a different routing table, <b>vrf-customer1.inet.0</b>. The second term does the opposite by using <b>rib vrf-customer1.inet.0</b> in the match condition and <b>vrf-customer2.inet.0</b> in the <b>next-table</b> action.</p> <pre> term 1 {   from {     protocol bgp;     rib vrf-customer2.inet.0;     community customer;   }   then {     next-hop next-table vrf-customer1.inet.0;   } } term 2 {   from {     protocol bgp;     rib vrf-customer1.inet.0;     community customer;   }   then {     next-hop next-table vrf-customer2.inet.0;   } } </pre> <p>If you specify <b>reject</b>, the next-hop address is replaced by a reject next hop.</p>
<b>origin value</b>	<p>(BGP only) Set the BGP origin attribute to one of the following values:</p> <ul style="list-style-type: none"> <li>• <b>igp</b>—Path information originated within the local AS.</li> <li>• <b>egp</b>—Path information originated in another AS.</li> <li>• <b>incomplete</b>—Path information learned by some other means.</li> </ul>

Table 9: Actions That Manipulate Route Characteristics (*continued*)

Action	Description
<b>preference</b> <i>preference</i> <b>preference2</b> <i>preference</i>	<p>Set the preference value. You can specify a primary preference value (<b>preference</b>) and a secondary preference value (<b>preference2</b>). The preference value can be a number in the range from 0 through 4,294,967,295 (<math>2^{32} - 1</math>). A lower number indicates a more preferred route.</p> <p>To specify even finer-grained preference values, see the <b>color</b> and <b>color2</b> actions in this table.</p> <p>If you set the preference with the <b>preference</b> action, the new preference remains associated with the route. The new preference is internal to the Junos OS and is not transitive.</p>
<b>preference</b> (add   subtract) <i>number</i> <b>preference2</b> (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 (<math>2^{32} - 1</math>), the value is set to <math>2^{32} - 1</math>. 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>
<b>priority</b> (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 <b>high</b> are installed first, while prefixes assigned a priority of <b>low</b> are installed last.</p> <p><b>NOTE:</b> An 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 <b>reject</b> terminating action for a nonexternal route, then the <b>reject</b> action is ignored and the route is accepted anyway.</p>
<b>source-class</b> <i>source-class-name</i>	<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"> <li>• Configure group source prefixes by configuring a routing policy; see <a href="#">“Defining Routing Policies” on page 50</a> and <a href="#">“Routing Policy Examples” on page 84</a>.</li> <li>• Apply that routing policy to the forwarding table with the corresponding source class; see <a href="#">“Applying Routing Policies to the Forwarding Table” on page 75</a>.</li> <li>• Enable packet counting on one or more interfaces by including the <b>source-class-usage</b> <i>interface-name</i> statement at the <b>[edit interfaces logical-unit-number unit family inet accounting]</b> hierarchy level. Also, follow the <b>source-class-usage</b> statement with the <b>input</b> or <b>output</b> 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 <b>[edit interfaces interface-name unit family inet accounting source-class-usage (input   output   input output) unit-number]</b>. See the example in <a href="#">“Routing Policy Examples” on page 84</a>.</li> <li>• View the output by using one of the following commands: <b>show interfaces interface-name source-class source-class-name</b>, <b>show interfaces interface-name extensive</b>, or <b>show interfaces interface-name statistics</b> (see the Junos OS Operational Mode Commands).</li> <li>• To configure a packet count based on the destination address, use the <b>destination-class</b> statement described in this table.</li> <li>• For a detailed source-class usage example configuration, see the Source Class Usage.</li> </ul>
<b>ssm-source</b> [ <i>addresses</i> ];	<p>Specify one or more IPv4 or IPv6 source addresses for the source-specific multicast (SSM) policy</p>



Table 9: Actions That Manipulate Route Characteristics (*continued*)

Action	Description
<b>ssm-source</b> [ <i>addresses</i> ];	Specify one or more IPv4 or IPv6 source addresses for the source-specific multicast (SSM) policy.
<b>tag</b> <i>tag</i> <b>tag2</b> <i>tag</i>	<p>Set the tag value. You can specify two tag strings: <b>tag</b> (for the first string) and <b>tag2</b> (a second string). These values are local to the router.</p> <ul style="list-style-type: none"> <li>For OSPF routes the <b>tag</b> action sets the 32-bit tag field in OSPF external link-state advertisement (LSA) packets.</li> <li>For IS-IS routes, the <b>tag</b> action sets the 32-bit flag in the IS-IS IP prefix type length values (TLV).</li> <li>For RIPv2 routes, the <b>tag</b> action sets the route-tag community. The <b>tag2</b> option is not supported.</li> </ul>
<b>tag</b> (add   subtract) <i>number</i> <b>tag2</b> (add   subtract) <i>number</i>	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.
<b>validation-state</b>	<p>When BGP origin validation is configured, set the validation state of a route prefix to valid, invalid, or unknown.</p> <p>The route validation database contains route origin authorization (ROA) records that map route prefixes to expected originating autonomous systems (ASs). This prevents the accidental advertisement of invalid routes.</p> <p>See Example: Configuring Origin Validation for BGP.</p>

## Configuring the Default Action in Routing Policies

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: Configuring the Default Action in a Routing 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;
      }
    }
  }
}

```

### Configuring a Final Action in Routing Policies

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;
}
}

```

## Logging Matches to a Routing Policy Term

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

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 is logged in the trace file. If a term specifies a trace action only, the word **<default>** is logged.

## Configuring Separate Actions for Routes in Route Lists

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 for Use in Routing Policy Match Conditions” on page 155](#).

## Applying Routing Policies and Policy Chains to Routing Protocols

---

For a routing policy to take effect, you must apply it to either a routing protocol or the forwarding table.

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. [“Protocols That Can Be Imported to and Exported from the Routing Table” on page 255](#) summarizes the import and export policy support for each routing protocol and the level at which you can apply these policies.

To apply one or more routing policies 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*]**

An ordered set of policies is referred to as a *policy chain*.

In the **import** statement, list the names of one or more 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 one or more 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 in a chain 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 information about how the policy framework software evaluates routing policies and policy chains, see [“Evaluating a Routing Policy” on page 26](#).

## Applying Policy Expressions to Routes Exported from Routing Tables

---

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 10 on page 71](#) summarizes the policy actions and their corresponding TRUE and FALSE values and flow control action values. [Table 11 on page 71](#) describes the logical operators. For complete information about policy expression evaluation, see [“How a Policy Expression Is Evaluated” on page 73](#).

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 10: 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 11: Policy Expression Logical Operators**

Logical Operator	Policy Expression Logic	How Logical Operator Affects Policy Expression Evaluation
<b>&amp;&amp;</b> (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>
<b>  </b> (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>

Table 11: Policy Expression Logical Operators (*continued*)

Logical Operator	Policy Expression Logic	How Logical Operator Affects Policy Expression Evaluation
! (Logical NOT)	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>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 72](#)
- [How a Policy Expression Is Evaluated on page 73](#)
- [Example: Evaluating Policy Expressions on page 74](#)

## 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 73](#).) 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 10 on page 71](#).)
2. The software takes the value of TRUE or FALSE and evaluates it against the logical operator used in the policy expression (see [Table 11 on page 71](#)). Based upon the logical operator used, the software determines whether or not to evaluate the next policy, if one is present.

The policy framework software uses a shortcut method of evaluation: if the result of evaluating a policy predetermines the value of the entire policy expression, the software does not evaluate the subsequent policies in the expression. For example, if the policy expression uses the logical AND operator and the evaluation of a policy returns the value of FALSE, the software does not evaluate subsequent policies in the expression because the final value of the expression is guaranteed to be FALSE no matter what the values of the unevaluated policies.

3. The software performs Step 1 and Step 2 for each subsequent routing policy in the policy expression, if they are present and it is necessary 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 10 on page 71](#).) 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 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.

## 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 26](#).

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 Per-Packet Load Balancing” on page 186](#).

**Related  
Documentation**

- Junos OS Class of Service Configuration Guide

---

## Configuring Dynamic Routing Policies

---

The verification process required to commit configuration changes can entail a significant amount of overhead and time. For example, changing a prefix in one line of a routing policy that is 20,000 lines long can take up to 20 seconds to commit. It can be useful to be able to commit routing policy changes much more quickly.

In Junos OS Release 9.5 and later, you can configure routing policies and certain routing policy objects in a dynamic database that is not subject to the same verification required in the standard configuration database. As a result, the time it takes to commit changes to the dynamic database is much shorter than for the standard configuration database. You can then reference these policies and policy objects in routing policies you configure in the standard database. BGP is the only protocol to which you can apply routing policies that reference policies and policy objects configured in the dynamic database. After you configure and commit a routing policy based on the objects configured in the dynamic database, you can quickly update any existing routing policy by making changes to the dynamic database configuration.



**CAUTION:** Because the Junos OS does not validate configuration changes to the dynamic database, when you use this feature, you should test and verify all configuration changes before committing them.

---

This section discusses the following topics:

- [Configuring Routing Policies and Policy Objects in the Dynamic Database on page 77](#)
- [Configuring Routing Policies Based on Dynamic Database Configuration on page 77](#)
- [Applying Dynamic Routing Policies to BGP on page 79](#)

- Preventing Reestablishment of BGP Peering Sessions After NSR Routing Engine Switchover on page 79
- Example: Configuring a BGP Export Policy That References a Dynamic Routing Policy on page 80

## Configuring Routing Policies and Policy Objects in the Dynamic Database

Junos OS Release 9.5 and later support a configuration database, the *dynamic database*, which can be edited in a similar way to the standard configuration database but which is not subject to the same verification process to commit configuration changes. As a result, the time it takes to commit a configuration change is much faster. The policies and policy objects defined in the dynamic database can then be referenced in routing policies configured in the standard configuration. The dynamic database is stored in the `/var/run/db/juniper.dyn` directory.

To configure the dynamic database, enter the **configure dynamic** command to enter the configuration mode for the dynamic database:

```
user@host> configure dynamic
Entering configuration mode
```

```
[edit dynamic]
user@host#
```

In this dynamic configuration database, you can configure the following statements at the **[edit policy-options]** hierarchy level:

- **as-path** *name*
- **as-path-group** *group-name*
- **community** *community-name*
- **condition** *condition-name*
- **prefix-list** *prefix-list-name*
- **policy-statement** *policy-statement-name*



**NOTE:** No other configuration is supported at the **[edit dynamic]** hierarchy level.

Use the **policy-statement** *policy-statement-name* statement to configure routing policies as you would in the standard configuration database.

To exit configuration mode for the dynamic database, issue the **exit configuration-mode** command from any level within the **[edit dynamic]** hierarchy, or use the **exit** command from the top level.

## Configuring Routing Policies Based on Dynamic Database Configuration

In the standard configuration mode, you can configure routing policies that reference policies and policy objects configured at the **[edit dynamic]** hierarchy level in the dynamic

database. To define a routing policy that references the dynamic database configuration, include the **dynamic-db** statement at the **[edit policy-options policy-statement *policy-statement-name*]** hierarchy level:

```
[edit policy-options]
policy-statement policy-statement-name {
  dynamic-db;
}
```

You can also define specific policy objects based on the configuration of these objects in the dynamic database. To define a policy object based on the dynamic database, include the **dynamic-db** statement with the following statements at the **[edit policy-options]** hierarchy level:

- **as-path** *name*
- **as-path-group** *group-name*
- **community** *community-name*
- **condition** *condition-name*
- **prefix-list** *prefix-list-name*

In the standard configuration, you can also define a routing policy that references any policy object you have configured in the standard configuration that references an object configured in the dynamic database.

For example, in standard configuration mode, you configure a prefix list **prefix-list pl2** that references a prefix list, also named **prefix-list pl2**, that has been configured in the dynamic database:

```
[edit policy-options]
prefix-list pl2 {
  dynamic-db; # Reference a prefix list configured in the dynamic database.
}
```

You then configure a routing policy in the standard configuration that includes **prefix-list pl2**:

```
[edit policy-options]
policy-statement one {
  term term1 {
    from {
      prefix-list pl2; # Include the prefix list configured in the standard configuration
                       # database, but which references a prefix list configured in the dynamic database.
    }
    then accept;
  }
  then reject;
}
```

If you need to update the configuration of **prefix-list pl2**, you do so in the dynamic database configuration using the **[edit dynamic]** hierarchy level. This enables you to make commit configuration changes to the prefix list more quickly than you can in the standard configuration database.



**NOTE:** If you are downgrading the Junos OS to Junos OS Release 9.4 or earlier, you must first delete any routing policies that reference the dynamic database. That is, you must delete any routing policies or policy objects configured with the `dynamic-db` statement.

## Applying Dynamic Routing Policies to BGP

BGP is the only routing protocol to which you can apply routing policies that reference the dynamic database configuration. You must apply these policies in the standard configuration. Dynamic policies can be applied to BGP export or import policy. They can also be applied at the global, group, or neighbor hierarchy level.

To apply a BGP export policy, include the `export [ policy-names ]` statement at the `[edit protocols bgp]`, `[edit protocols bgp group group-name]`, or `[edit protocols bgp group group-name neighbor address]` hierarchy level.

```
[edit]
protocols
  bgp {
    export [ policy-names ];
  }
}
```

To apply a BGP import policy, include the `import [ policy-names ]` statement at the `[edit protocols bgp]`, `[edit protocols bgp group group-name]`, or `[edit protocols bgp group group-name neighbor address]` hierarchy level.

```
[edit]
protocols
  bgp {
    import [ policy-names ];
  }
}
```

Include one or more policy names configured in that standard configuration at the `[edit policy-options policy-statement]` hierarchy level that reference policies configured in the dynamic database.

## Preventing Reestablishment of BGP Peering Sessions After NSR Routing Engine Switchover

If you have active nonstop routing (NSR) enabled, the dynamic database is not synchronized with the backup Routing Engine. As a result, if a switchover to a backup Routing Engine occurs, import and export policies running on the master Routing Engine at the time of the switchover might no longer be available. Therefore, you might want to prevent a BGP peering session from automatically being reestablished as soon as a switchover occurs.

You can configure the router not to reestablish a BGP peering session after an active nonstop routing switchover either for a specified period or until you manually reestablish the session. Include the `idle-after-switchover (seconds | forever)` statement at the `[edit`

`protocols bgp`], `[edit protocols bgp group group-name]`, or `[edit protocols bgp group group-name neighbor address]` hierarchy level:

```
[edit]
bgp {
  protocols {
    idle-after-switch-over (seconds | never);
  }
}
```

For ***seconds***, specify a value from 1 through 4,294,967,295 ( $2^{32} - 1$ ). The BGP peering session is not reestablished until after the specified period. If you specify the **forever** option, the BGP peering session is not established until you issue the **clear bgp neighbor** command. For more information about configuring active nonstop routing, see the Junos OS High Availability Configuration Guide.

### Example: Configuring a BGP Export Policy That References a Dynamic Routing Policy

In this example, you configure the following in the dynamic database: a routing policy, **policy-statement one**, and two prefix lists: **prefix-list pl1**, which is referenced in **policy-statement one**, and **prefix-list pl2**.

In the standard configuration database, you configure the following routing policies and policy objects: a routing policy (**policy-statement one**) that references the routing policy with the same name configured in the dynamic database; a prefix list (**prefix-list pl2**) that references the prefix list with the same name configured in the dynamic database; and a routing policy (**policy-statement two**) that references the prefix list **prefix-list pl2** you configured in the standard configuration and references the prefix list with the same name configured in the dynamic database.

You then create and apply a policy expression that includes **policy-statement one** and **policy-statement two** to BGP export policy in the standard configuration.

In the dynamic database, configure **policy-statement one**, **prefix-list pl1**, and **prefix-list pl2**:

```
[edit dynamic]
policy-options {
  prefix-list pl1 {
    8.8.0.0/16;
    12.12.12.3/32
  }
  prefix-list pl2;
  10.10.0.0/16
}
policy-statement one {
  term term1
  from {
    prefix-list pl1
  }
  then accept;
}
then reject;
}
```

In the standard configuration database, configure **policy-statement one** based on the policy with the same name configured in the dynamic database. In addition, configure **prefix-list pl2**, which references the prefix list with the same name in the dynamic database, and configure **policy-statement two**, which references **prefix-list pl2**.

```
[edit]
policy-options {
  prefix-list pl2;
  dynamic-db; # Reference 'prefix-list pl2' configured in the dynamic database.
}
policy-statement two {
  term term1 {
    from {
      prefix-list pl2; # Configure a routing policy that includes a prefix list that
                      # references the dynamic database.
    }
    then accept;
  }
  then reject;
}
policy-statement one;
dynamic-db; # Configure a policy in the standard configuration that references
           # the policy configured in the dynamic database.
}
```

Apply a policy expression that includes routing policies **one** and **two** to BGP export policy for an internal BGP group **test**.

```
[edit]
protocols {
  bgp {
    group test;
    type internal;
    local-address 10.255.245.44;
    export ( one && two ); # If routing policy one returns the value of TRUE, policy # two
                          # is evaluated.
                          # If policy one returns the value of FALSE, the evaluation of the expression # ends
                          # and policy two is not evaluated.
  }
}
```

## Forwarding Packets to 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.

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, you can issue the **test policy** command to ensure that the policy produces the results that you expect:

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



## 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
```

## Example: Creating a Routing Policy

This example shows how to create a simple routing policy.

- [Requirements on page 83](#)
- [Overview on page 83](#)
- [Configuration on page 83](#)
- [Verification on page 84](#)

### Requirements

Before you begin, determine what you want to accomplish with the policy, configure router interfaces, and configure routing protocols, as explained in “[Routing Policies Configuration Overview](#)” on page 10.

### Overview

In this example, you create a routing policy called policy1.

### Configuration

#### Step-by-Step Procedure

To create a routing policy:

1. Create the routing policy.

```
[edit]
user@host# set policy-options policy-statement policy1
```

2. Commit the configuration if you are done configuring the device.

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



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

---

## Verification

To verify your configuration, use the **show policy-options** command.

### Related Documentation

- [Routing Policies Configuration Overview on page 10](#)
- [Importing and Exporting Routes on page 17](#)

## Routing Policy Examples

---

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

- [Example: Defining a Routing Policy from BGP to IS-IS on page 97](#)
- [Example: Using Routing Policy to Set a Preference on page 98](#)
- [Example: Importing and Exporting Access and Access-Internal Routes in a Routing Policy on page 98](#)
- [Example: Exporting Routes to IS-IS on page 99](#)
- [Example: Applying Export and Import Policies to BGP Peer Groups on page 99](#)
- [Example: Defining a Routing Policy Based on the Number of BGP Communities on page 100](#)
- [Example: Applying a Prefix to Routes Learned from a Peer on page 100](#)
- [Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS on page 100](#)
- [Example: Redistributing OSPF Routes into BGP on page 101](#)
- [Example: Exporting Direct Routes Into IS-IS on page 102](#)
- [Example: Exporting Internal IS-IS Level 1 Routes to Level 2 on page 102](#)
- [Example: Exporting IS-IS Level 2 Routes to Level 1 on page 102](#)
- [Example: Assigning Different Forwarding Next-Hop LSPs to Different Destination Prefixes on page 103](#)
- [Example: Grouping Destination Prefixes on page 104](#)
- [Example: Grouping Source Prefixes on page 105](#)
- [Example: Grouping Source and Destination Prefixes into a Forwarding Class on page 106](#)

- [Example: Accepting Routes with Specific Destination Prefixes on page 108](#)
- [Example: Accepting Routes from BGP with a Specific Destination Prefix on page 108](#)

## Example: Creating a Routing Policy Term

This example shows how to create a routing policy term.

- [Requirements on page 85](#)
- [Overview on page 85](#)
- [Configuration on page 85](#)
- [Verification on page 85](#)

### Requirements

Before you begin, determine what you want to accomplish with the policy, configure router interfaces, and configure routing protocols, as explained in “[Routing Policies Configuration Overview](#)” on page 10.

### Overview

In this example, you create a routing policy called `policy1` and a term for the policy called `term1`.

### Configuration

#### Step-by-Step Procedure

To configure a routing policy term:

1. Create the routing policy.  
  
[edit]  
user@host# **edit policy-options policy-statement policy1**
2. Create the policy term.  
  
[edit policy-options policy-statement policy1]  
user@host# **set term term1**
3. Commit the configuration if you are done configuring the device.  
  
[edit policy-options policy-statement policy1]  
user@host# **commit**



**NOTE:** The policy does not take effect until you apply it.

### Verification

To verify your configuration, use the **show policy-options** command.

#### Related Documentation

- [Junos OS Feature Support Reference for SRX Series and J Series Devices](#)

- [Routing Policies Configuration Overview on page 10](#)
- [Routing Policy Terms on page 23](#)

## Example: Rejecting Known Invalid Routes

---

This example shows how to create route-based match conditions for a routing policy.

- [Requirements on page 86](#)
- [Overview on page 86](#)
- [Configuration on page 86](#)
- [Verification on page 87](#)

### Requirements

Before you begin, configure router interfaces and configure routing protocols, as explained in [“Routing Policies Configuration Overview” on page 10](#).

### Overview

In this example, you create a policy called `rejectpolicy1` that rejects routes with a mask of /8 and greater (/8, /9, /10, and so on) that have the first 8 bits set to 0. This policy also accepts routes less than 8 bits in length by creating a mask of 0/0 up to /7.

### Configuration

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

```
set policy-options policy-statement rejectpolicy1 term rejectterm1 from route-filter
  0.0.0.0/0 upto /7 accept
set policy-options policy-statement rejectpolicy1 term rejectterm1 from route-filter
  0.0.0.0/8 orlonger reject
set policy-options policy-statement test term 1 from protocol direct
```

**Step-by-Step Procedure** To create a policy that rejects known invalid routes:

1. Create the routing policy.  

```
[edit]
user@host# edit policy-options policy-statement rejectpolicy1
```
2. Create the policy term.  

```
[edit policy-options policy-statement rejectpolicy1]
user@host# edit term rejectterm1
```
3. Create a mask that specifies which routes to accept.  

```
[edit policy-options policy-statement rejectpolicy1 term rejectterm1]
user@host# set from route-filter 0/0 upto /7 accept
```

4. Create a mask that specifies which routes to reject.

```
[edit policy-options policy-statement rejectpolicy1 term rejectterm1]
user@host# set from route-filter 0/8 orlonger reject
```

**Results** Confirm your configuration by entering the **show policy-options** command from configuration mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show policy-options
policy-statement rejectpolicy1 {
  term rejectterm1 {
    from {
      route-filter 0.0.0.0/0 upto /7 accept;
      route-filter 0.0.0.0/8 orlonger reject;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying the Route-Based Match Conditions on page 87](#)

### Verifying the Route-Based Match Conditions

**Purpose** Verify that the policy and term are configured on the device with the appropriate route-based match conditions.

**Action** From operational mode, enter the **show policy-options** command.

**Related Documentation**

- [Junos OS Feature Support Reference for SRX Series and J Series Devices](#)
- [Routing Policies Configuration Overview on page 10](#)
- [Route-Based Match Conditions on page 21](#)
- [Example: Grouping Source and Destination Prefixes into a Forwarding Class on page 106](#)

## Examples: Configuring Static Routes

- [Understanding Basic Static Routing on page 88](#)
- [Example: Configuring a Basic Set of Static Routes on page 88](#)
- [Example: Configuring IPv6 Static Routes on page 93](#)

## Understanding Basic Static Routing

Routes that are permanent fixtures in the routing and forwarding tables are often configured as static routes. These routes generally do not change, and often include only one or very few paths to the destination.

To create a static route in the routing table, you must, at minimum, define the route as static and associate a next-hop address with it. The static route in the routing table is inserted into the forwarding table when the next-hop address is reachable. All traffic destined for the static route is transmitted to the next-hop address for transit.

You can specify options that define additional information about static routes that is included with the route when it is installed in the routing table. All static options are optional.

### Example: Configuring a Basic Set of Static Routes

This example shows how to configure a basic set of static routes.

- [Requirements on page 88](#)
- [Overview on page 88](#)
- [Configuration on page 89](#)
- [Verification on page 91](#)

---

#### Requirements

In this example, no special configuration beyond device initialization is required.

---

#### Overview


There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.

In this example, you configure the static route 192.168.47.0/24 from the provider network to the customer network, using the next-hop address of 172.16.1.2. You also configure a static default route of 0.0.0.0/0 from the customer network to the provider network, using a next-hop address of 172.16.1.1.

For demonstration purposes, some loopback interfaces are configured on Device B and Device D. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

[Figure 12 on page 89](#) shows the sample network.

Figure 12: Customer Routes Connected to a Service Provider



### Configuration

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

**Device B**

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
```

**Device D**

```
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
```

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

To configure basic static routes:

1. On Device B, configure the interfaces.

```
[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24
user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32
```

2. On Device B, create a static route and set the next-hop address.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2
```

3. If you are done configuring Device B, commit the configuration.

```
[edit interfaces]
user@B# commit
```

4. On Device D, configure the interfaces.

```
[edit]
user@D# set ge-1/2/0 unit 1 description D->B
user@D# set ge-1/2/0 unit 1 family inet address 172.16.1.2/24
user@D# set lo0 unit 2 family inet address 192.168.47.5/32
user@D# set lo0 unit 2 family inet address 192.168.47.6/32
```

5. On Device D, create a static route and set the next-hop address.

```
[edit routing-options]
user@D# set static route 0.0.0.0/0 next-hop 172.16.1.1
```

6. If you are done configuring Device D, commit the configuration.

```
[edit]
user@D# commit
```

### **Results**

Confirm your configuration by issuing the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device B user@B# show interfaces
ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
  }
}
lo0 {
  unit 57 {
    family inet {
      address 10.0.0.1/32;
      address 10.0.0.2/32;
    }
  }
}

user@B# show routing-options
static {
  route 192.168.47.0/24 next-hop 172.16.1.2;
}

Device D user@D# show interfaces
ge-1/2/0 {
```



```
unit 1 {
  description D->B;
  family inet {
    address 172.16.1.2/24;
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.47.5/32;
      address 192.168.47.6/32;
    }
  }
}

user@D# show routing-options
static {
  route 0.0.0.0/0 next-hop 172.16.1.1;
}
```

---

### Verification

Confirm that the configuration is working properly.

- [Checking the Routing Tables on page 91](#)
- [Pinging the Remote Addresses on page 92](#)

#### ***Checking the Routing Tables***

**Purpose** Make sure that the static routes appear in the routing tables of Device B and Device D.

**Action** user@B> show route  
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
10.0.0.1/32      *[Direct/0] 00:29:43
                  > via lo0.57
10.0.0.2/32      *[Direct/0] 00:29:43
                  > via lo0.57
172.16.1.0/24    *[Direct/0] 00:34:40
                  > via ge-1/2/0.0
172.16.1.1/32    *[Local/0] 00:34:40
                  Local via ge-1/2/0.0
192.168.47.0/24  *[Static/5] 00:31:23
                  > to 172.16.1.2 via ge-1/2/0.0
```

user@D> show route  
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
0.0.0.0/0        *[Static/5] 00:31:24
                  > to 172.16.1.1 via ge-1/2/0.1
172.16.1.0/24    *[Direct/0] 00:35:21
                  > via ge-1/2/0.1
172.16.1.2/32    *[Local/0] 00:35:21
                  Local via ge-1/2/0.1
192.168.47.5/32  *[Direct/0] 00:35:22
                  > via lo0.2
192.168.47.6/32  *[Direct/0] 00:35:21
                  > via lo0.2
```

**Meaning** The static routes are in the routing tables.

#### *Pinging the Remote Addresses*

**Purpose** Verify that the static routes are working.

From Device B, ping one of the loopback interface addresses on Device D.

From Device D, ping one of the loopback interface addresses on Device B.

**Action** user@B> ping 192.168.47.5  
PING 192.168.47.5 (192.168.47.5): 56 data bytes  
64 bytes from 192.168.47.5: icmp\_seq=0 ttl=64 time=156.126 ms  
64 bytes from 192.168.47.5: icmp\_seq=1 ttl=64 time=120.393 ms  
64 bytes from 192.168.47.5: icmp\_seq=2 ttl=64 time=175.361 ms

user@D> ping 10.0.0.1  
PING 10.0.0.1 (10.0.0.1): 56 data bytes  
64 bytes from 10.0.0.1: icmp\_seq=0 ttl=64 time=1.315 ms  
64 bytes from 10.0.0.1: icmp\_seq=1 ttl=64 time=31.819 ms  
64 bytes from 10.0.0.1: icmp\_seq=2 ttl=64 time=1.268 ms

## Example: Configuring IPv6 Static Routes

This example shows how to configure static routes when the interfaces have IPv6 addresses.

- [Requirements on page 93](#)
- [Overview on page 93](#)
- [Configuration on page 94](#)
- [Verification on page 96](#)

### Requirements

---

In this example, no special configuration beyond device initialization is required.

### Overview

---


There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.

In this example, you configure a static default route of `::/0`, using a next-hop address `2001:db8:0:1:2a0:a502:0:1da`.

For demonstration purposes, some loopback interfaces are configured on Device A and Device E. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

[Figure 13 on page 94](#) shows the sample network.

Figure 13: Customer Routes Connected to a Service Provider



### Configuration

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

**Device A**

```
set interfaces fe-1/2/0 unit 1 description to-E
set interfaces fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1:2a0:a502:0:1da/64
set interfaces lo0 unit 1 family inet6 address 2001:db8::1/128 primary
set interfaces lo0 unit 1 family inet6 address 2001:db8::2/128
set interfaces lo0 unit 1 family inet6 address 2001:db8::3/128
set routing-options rib inet6.0 static route 2001:db8::5/128 next-hop
  2001:db8:0:1:2a0:a502:0:19da
```

**Device E**

```
set interfaces fe-1/2/0 unit 25 description to-A
set interfaces fe-1/2/0 unit 25 family inet6 address 2001:db8:0:1:2a0:a502:0:19da/64
set interfaces lo0 unit 5 family inet6 address 2001:db8::5/128
set routing-options rib inet6.0 static route ::/0 next-hop 2001:db8:0:1:2a0:a502:0:1da
```

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

To configure basic static routes:

- On Device A, configure the interfaces.
 

```
[edit interfaces]
set interfaces fe-1/2/0 unit 1 description to-E
set interfaces fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1:2a0:a502:0:1da/64

set interfaces lo0 unit 1 family inet6 address 2001:db8::1/128 primary
```

```
set interfaces lo0 unit 1 family inet6 address 2001:db8::2/128
set interfaces lo0 unit 1 family inet6 address 2001:db8::3/128
```

2. On Device A, create a static route to Device E's loopback address and set the next-hop address.

This ensures that Device A has a route back to Device E.

```
[edit routing-options]
set rib inet6.0 static route 2001:db8::5/128 next-hop 2001:db8:0:1:2a0:a502:0:19da
```

3. If you are done configuring Device A, commit the configuration.

```
[edit interfaces]
user@A# commit
```

4. On Device E, configure the interfaces.

```
[edit]
set interfaces fe-1/2/0 unit 25 description to-A
set interfaces fe-1/2/0 unit 25 family inet6 address 2001:db8:0:1:2a0:a502:0:19da/64
set interfaces lo0 unit 5 family inet6 address 2001:db8::5/128
```

5. On Device E, create a static default route and set the next-hop address.

```
[edit routing-options]
set routing-options rib inet6.0 static route ::/0 next-hop 2001:db8:0:1:2a0:a502:0:1da
```

6. If you are done configuring Device E, commit the configuration.

```
[edit]
user@E# commit
```

### Results

Confirm your configuration by issuing the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device A user@A# show interfaces
fe-1/2/0 {
  unit 1 {
    description to-E;
    family inet6 {
      address 2001:db8:0:1:2a0:a502:0:1da/64;
    }
  }
}
lo0 {
  unit 1 {
    family inet6 {
      address 2001:db8::1/128 {
        primary;
      }
      address 2001:db8::2/128;
      address 2001:db8::3/128;
    }
  }
}
```

```
user@A# show routing-options
rib inet6.0 {
  static {
    route 2001:db8::5/128 next-hop 2001:db8:0:1:2a0:a502:0:19da;
  }
}

Device E user@E# show interfaces
fe-1/2/0 {
  unit 25 {
    description to-A;
    family inet6 {
      address 2001:db8:0:1:2a0:a502:0:19da/64;
    }
  }
}
lo0 {
  unit 5 {
    family inet6 {
      address 2001:db8::5/128;
    }
  }
}

user@E# show routing-options
rib inet6.0 {
  static {
    route ::/0 next-hop 2001:db8:0:1:2a0:a502:0:1da;
  }
}
```

---

### Verification

Confirm that the configuration is working properly.

- [Checking the Routing Tables on page 96](#)
- [Pinging the Remote Addresses on page 97](#)

#### *Checking the Routing Tables*

**Purpose** Make sure that the static routes appear in the routing tables of Device A and Device E.

```

Action user@A> show route protocol static
inet6.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::5/128    *[Static/5] 00:27:46
                  > to 2001:db8:0:1:2a0:a502:0:19da via fe-1/2/0.1

user@E> show route protocol static
inet6.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0              *[Static/5] 00:19:11
                  > to 2001:db8:0:1:2a0:a502:0:1da via fe-1/2/0.25

```

**Meaning** The static routes are in the routing tables.

### *Pinging the Remote Addresses*

**Purpose** Verify that the static routes are working.

From Device A, ping one of the loopback interface addresses on Device E.

From Device E, ping one of the loopback interface addresses on Device A.

```

Action user@A> ping 2001:db8::5
PING6(56=40+8+8 bytes) 2001:db8:0:1:2a0:a502:0:1da --> 2001:db8::5
16 bytes from 2001:db8::5, icmp_seq=0 hlim=64 time=1.790 ms
16 bytes from 2001:db8::5, icmp_seq=1 hlim=64 time=1.529 ms
16 bytes from 2001:db8::5, icmp_seq=2 hlim=64 time=1.531 ms

user@E> ping 2001:db8::3
PING6(56=40+8+8 bytes) 2001:db8:0:1:2a0:a502:0:19da --> 2001:db8::3
16 bytes from 2001:db8::3, icmp_seq=0 hlim=64 time=2.146 ms
16 bytes from 2001:db8::3, icmp_seq=1 hlim=64 time=1.964 ms
16 bytes from 2001:db8::3, icmp_seq=2 hlim=64 time=1.550 ms

```

## Example: 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 18.](#)) 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;
      }
    }
  }
}

```

```
    }  
  }  
}
```

---

## Example: 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 18.](#))

```
[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;  
      }  
    }  
  }  
}
```

---

## Example: 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 Junos OS Subscriber Management, Release 12.3.)

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;
    }
  }
}
```

### Example: Exporting Routes to IS-IS

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 ];
  }
}
```

### Example: Applying Export and Import Policies to BGP Peer Groups

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;
    }
  }
}
```

## Example: Defining a Routing Policy Based on the Number of BGP Communities

Create a policy that accepts BGP routes based on the number of BGP communities. In this example, routes can contain two, three, or four communities to be considered a match. For example, if a route contains three communities, it is considered a match and is accepted. If a route contains one community, it is not considered a match and is rejected.



**NOTE:** The `community-count` attribute only works with standard communities. It does not work with extended communities.

```
[edit]
policy-options {
  policy-statement import-bgp {
    term community {
      from {
        community-count 2 orhigher;
        community-count 4 orlower;
      }
      then {
        accept;
      }
    }
  }
}
```

## Example: Applying a Prefix to Routes Learned from a Peer

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 ];
      }
    }
  }
}
```

## 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 into a Forwarding Class

---

This example shows how to group source and destination prefixes into a forwarding class.

- [Requirements on page 106](#)
- [Overview on page 106](#)
- [Configuration on page 106](#)
- [Verification on page 108](#)

### Requirements

Before you begin, configure router interfaces and configure routing protocols, as explained in [“Routing Policies Configuration Overview” on page 10](#).

### Overview

In this example, you configure a routing policy called `policy1` and a corresponding routing term called `term1`. Within the term, you configure the route filter to include source routes greater than or equal to `10.210.0.0/16` and destination routes greater than or equal to `10.215.0.0/16`. Then you group the source and destination prefixes into a forwarding class called `forwarding-class1` and apply `policy1` to the forwarding table. The routing policy is evaluated when routes are being exported from the routing table into the forwarding table. Only the active routes are exported from the routing table.

### Configuration

#### CLI Quick Configuration

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

```
set policy-options policy-statement policy1 term term1 from route-filter 10.210.0.0/16
  orlonger
set policy-options policy-statement policy1 term term1 from route-filter 10.215.0.0/16
  orlonger
set policy-options policy-statement policy1 term term1 then forwarding-class
  forwarding-class1
set routing-options forwarding-table export policy1
```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see [Using the CLI Editor in Configuration Mode](#) in the CLI User Guide.

To group source and destination prefixes in a forwarding class:

1. Create the routing policy.  
  
[edit]  
user@host# **edit policy-options policy-statement policy1**



2. Create the routing term.

```
[edit policy-options policy-statement policy1]
user@host# edit term term1
```

3. Specify the source routes to include in the route filter.

```
[edit policy-options policy-statement policy1 term term1]
user@host# set from route-filter 10.210.0.0/16 orlonger
```

4. Specify the destination routes to include in the route filter.

```
[edit policy-options policy-statement policy1 term term1]
user@host# set from route-filter 10.215.0.0/16 orlonger
```

5. Group the source and destination prefixes into the forwarding class.

```
[edit policy-options policy-statement policy1 term term1]
user@host# set then forwarding-class forwarding-class1
```

6. Apply the routing policy to the forwarding table.

```
[edit]
user@host# set routing-options forwarding-table export policy1
```



**NOTE:** You can refer to the same routing policy one or more times in the same or different export statement.

**Results** Confirm your configuration by entering the **show policy-options** and **show routing-options** commands from configuration mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show policy-options
policy-statement policy1 {
  term term1 {
    from {
      route-filter 10.210.0.0/16 orlonger;
      route-filter 10.215.0.0/16 orlonger;
    }
    then forwarding-class forwarding-class1;
  }
}

user@host# show routing-options
forwarding-table {
  export policy1;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying the Forwarding Class on page 108](#)
- [Verifying the Routing Policy on page 108](#)

### Verifying the Forwarding Class

**Purpose** Verify that the forwarding table is applied to the routing policy.

**Action** From operational mode, enter the **show routing-options** command.

### Verifying the Routing Policy

**Purpose** Verify that the policy and term are configured on the device with the appropriate routes included in the forwarding class.

**Action** From operational mode, enter the **show policy-options** command.

**Related Documentation**

- [Routing Policies Configuration Overview on page 10](#)
- [Route-Based Match Conditions on page 21](#)
- [Example: Rejecting Known Invalid Routes on page 86](#)

## 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: Using Routing Policy in an ISP Network

---

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 14 on page 110](#)), 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 BGP.

Figure 14: 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 14 on page 110](#), 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 111](#)
- [Requesting Specific Routes on the Customer 2 Router on page 112](#)
- [Configuring a Peer Policy on ISP Router 3 on page 114](#)
- [Configuring Private and Exchange Peers on ISP Router 1 and 2 on page 116](#)
- [Configuring Locally Defined Static Routes on the Exchange Peer 2 Router on page 119](#)
- [Configuring Outbound and Generated Routes on the Private Peer 2 Router on page 119](#)

## 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 only for 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";
}

```

## Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Switches

- [Understanding the Default BGP Routing Policy on Packet Transport Switches on page 121](#)
- [Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Switches on page 123](#)

### Understanding the Default BGP Routing Policy on Packet Transport Switches

On PTX Series Packet Transport Switches, the default BGP routing policy differs from that of other Junos OS routing devices.

The PTX Series switches are MPLS transit platforms that do IP forwarding, typically using interior gateway protocol (IGP) routes. The PTX Series Packet Forwarding Engine (PFE) can accommodate a relatively small number of variable-length prefixes.



**NOTE:** A PTX Series switch can support full BGP routes in the control plane so that it can be used as a route reflector (RR). It can do exact-length lookup multicast forwarding and can build the multicast forwarding plane for use by the unicast control plane (for example, to perform a reverse-path forwarding lookup for multicast).

Given the PFE limitation, the default routing policy for PTX Series switches is for BGP routes not to be installed in the forwarding table. You can override the default routing policy and select certain BGP routes to install in the forwarding table.

The default behavior for load balancing and BGP routes on PTX Series switches is as follows. It has the following desirable characteristics:

- Allows you to override the default behavior without needing to alter the default policy directly
- Reduces the chance of accidental changes that nullify the defaults
- Sets no flow-control actions, such as accept and reject

The default routing policy on the PTX Series switches is as follows:

```
user@host# show policy-options | display inheritance defaults no-comments
policy-options {
  policy-statement junos-ptx-series-default {
    term t1 {
      from {
        protocol bgp;
        rib inet.0;
      }
      then no-install-to-fib;
    }
    term t2 {
      from {
        protocol bgp;
        rib inet6.0;
      }
      then no-install-to-fib;
    }
    term t3 {
      then load-balance per-packet;
    }
  }
}
routing-options {
  forwarding-table {
    default-export junos-ptx-series-default;
  }
}
user@host# show routing-options forwarding-table default-export | display inheritance
defaults no-comments
default-export junos-ptx-series-default;
```

As shown here, the **junos-ptx-series-default** policy is defined in **[edit policy-options]**. The policy is applied in **[edit routing-options forwarding-table]**, using the **default-export** statement. You can view these default configurations by using the **| display inheritance** flag.

Also, you can use the **show policy** command to view the default policy.

```
user@host> show policy junos-ptx-series-default
```



```

Policy junos-ptx-series-default:
  Term t1:
    from proto BGP
    inet.0
    then install-to-fib no
  Term t2:
    from proto BGP
    inet6.0
    then install-to-fib no
  Term t3:
    then load-balance per-packet

```



**CAUTION:** We strongly recommend that you do not alter the `junos-ptx-series-default` routing policy directly.

Junos OS chains the `junos-ptx-series-default` policy and any user-configured export policy. Because the `junos-ptx-series-default` policy does not use flow-control actions, any export policy that you configure is executed (by way of the implicit `next-policy` action) for every route. Thus you can override any actions set by the `junos-ptx-series-default` policy. If you do not configure an export policy, the actions set by `junos-ptx-series-default` policy are the only actions.

You can use the policy action `install-to-fib` to override the `no-install-to-fib` action.

Similarly, you can set the `load-balance per-prefix` action to override the `load-balance per-packet` action.

## Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Switches

This example shows how to override the default routing policy on packet transport switches, such as the PTX Series Packet Transport Switches.

- [Requirements on page 123](#)
- [Overview on page 123](#)
- [Configuration on page 124](#)
- [Verification on page 125](#)

### Requirements

This example requires Junos OS Release 12.1 or later.

### Overview

By default, the PTX Series switches do not install BGP routes in the forwarding table.

For PTX Series switches, the configuration of the `from protocols bgp` condition with the `then accept` action does not have the usual result that it has on other Junos OS routing devices. With the following routing policy on PTX Series switches, BGP routes do not get installed in the forwarding table.

```

user@host# show policy-options
policy-statement accept-no-install {

```

```

term 1 {
    from protocol bgp;
    then accept;
}
}
user@host# show routing-options
forwarding-table {
    export accept-no-install
}

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif
default              perm   0                rjct   36    2

```

No BGP routes are installed in the forwarding table. This is the expected behavior.

This example shows how to use the **then install-to-fib** action to effectively override the default BGP routing policy.

### Configuration

#### CLI Quick Configuration

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

```

set policy-options prefix-list install-bgp 66.0.0.1/32
set policy-options policy-statement override-ptx-series-default term 1 from prefix-list
install-bgp
set policy-options policy-statement override-ptx-series-default term 1 then load-balance
per-prefix
set policy-options policy-statement override-ptx-series-default term 1 then install-to-fib
set routing-options forwarding-table export override-ptx-series-default

```

#### Installing Selected BGP Routes in the Forwarding Table

#### Step-by-Step Procedure

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

To install selected BGP routes in the forwarding table:

1. Configure a list of prefixes to install in the forwarding table.
 

```

[edit policy-options prefix-list install-bgp]
user@host# set 66.0.0.1/32

```
2. Configure the routing policy, applying the prefix list as a condition.
 

```

[edit policy-options policy-statement override-ptx-series-default term 1]
user@host# set from prefix-list install-bgp
user@host# set then install-to-fib
user@host# set then load-balance per-prefix

```
3. Apply the routing policy to the forwarding table.

```
[edit routing-options forwarding-table]
user@host# set export override-ptx-series-default
```

**Results** From configuration mode, confirm your configuration by entering the **show policy-options** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
prefix-list install-bgp {
  66.0.0.1/32;
}
policy-statement override-ptx-series-default {
  term 1 {
    from {
      prefix-list install-bgp;
    }
    then {
      load-balance per-prefix;
      install-to-fib;
    }
  }
}

user@host# show routing-options
forwarding-table {
  export override-ptx-series-default;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

#### *Verifying That the Selected Route Is Installed in the Forwarding Table*

**Purpose** Make sure that the configured policy overrides the default policy

**Action** From operational mode, enter the **show route forwarding-table** command.

```
user@host> show route forwarding-table destination 66.0.0.1
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif
66.0.0.1/32          user    0                indr 2097159 3
                    ucst 2097156 2
                    5.1.0.2    ucst 574 1 et-6/0/0.1
                    5.2.0.2    ucst 575 1 et-6/0/0.2
```

**Meaning** This output shows that the route to 66.0.0.1/32 is installed in the forwarding table.

**Related Documentation**

- [Default Routing Policies and Actions on page 18](#)
- [Default Import and Export Policies for Protocols on page 257](#)
- [Examples: Configuring BGP Multipath](#)

## Understanding ECMP Flow-Based Forwarding

---

An equal-cost multipath (ECMP) set is formed when the routing table contains multiple next-hop addresses for the same destination with equal cost. (Routes of equal cost have the same preference and metric values.) If there is an ECMP set for the active route, the Junos OS software uses a hash algorithm to choose *one* of the next-hop addresses in the ECMP set to install in the forwarding table.

You can configure Junos OS so that multiple next-hop entries in an ECMP set are installed in the forwarding table. On Juniper Networks devices, per-flow load balancing can be performed to spread traffic across multiple paths between routing devices. On Juniper Networks security devices, source and destination IP addresses and protocols are examined to determine individual traffic flows. Packets for the same flow are forwarded on the same interface; the interface does not change when there are additions or changes to the ECMP set. This is important for features such as source NAT, where the translation is performed only during the first path of session establishment; IDP; ALG; and route-based VPN tunnels. If a packet arrives on a given interface in an ECMP set, the security device ensures that reverse traffic is forwarded through the same interface.



**NOTE:** ECMP flow-based forwarding on security devices applies only to IPv4 unicast traffic flows. Multicast and IPv6 flows are not supported.

---

On Juniper Networks security devices, the maximum number of next-hop addresses in an ECMP set that can be installed in the forwarding table is 16. If there are more than 16 next-hop addresses in an ECMP set, only the first 16 addresses are used.

In a chassis cluster deployment, a *local* interface is an interface that is on the same node as the interface on which a packet arrives, and a *remote* interface is an interface that is on the other chassis cluster node. If an ECMP route has both local and remote interfaces in a chassis cluster, then the local interface is favored for the next hop.

If a next-hop address is no longer part of the ECMP set or if it is removed from the routing table because of a route change, a flow that uses the next hop is rerouted and the session is not affected. Rerouting of the flow also occurs if there is a configuration change that takes away the next-hop address or if an administrator takes down the next-hop interface without deleting it. If a next-hop address is removed from the routing table because the interface is deleted or the session is intentionally cleared, the session is killed without being rerouted.



**NOTE:** We recommend that interfaces in an ECMP set be in the same security zone. If a flow is rerouted and the rerouted flow uses an interface in a different security zone than the original route, the session is killed.

---

To configure ECMP flow-based forwarding on Juniper Networks security devices, first define a load-balancing routing policy by including one or more **policy-statement** configuration statements at the **[edit policy-options]** hierarchy level, with the action

**load-balance per-packet**. Then apply the routing policy to routes exported from the routing table to the forwarding table. To do this, include the **forwarding-table** and **export** configuration statements at the **[edit routing-options]** hierarchy level.

**Related  
Documentation**

- [Example: Configuring ECMP Flow-Based Forwarding on page 127](#)
- [Understanding Routing Policies on page 15](#)
- [Routing Policy Match Conditions on page 258](#)
- [Default Policies and Actions on page 8](#)

---

## Example: Configuring ECMP Flow-Based Forwarding

---

This example shows how to configure ECMP flow-based forwarding.

- [Requirements on page 127](#)
- [Overview on page 127](#)
- [Configuration on page 128](#)
- [Verification on page 131](#)

### Requirements

No special configuration beyond device initialization is required before configuring this feature.

### Overview

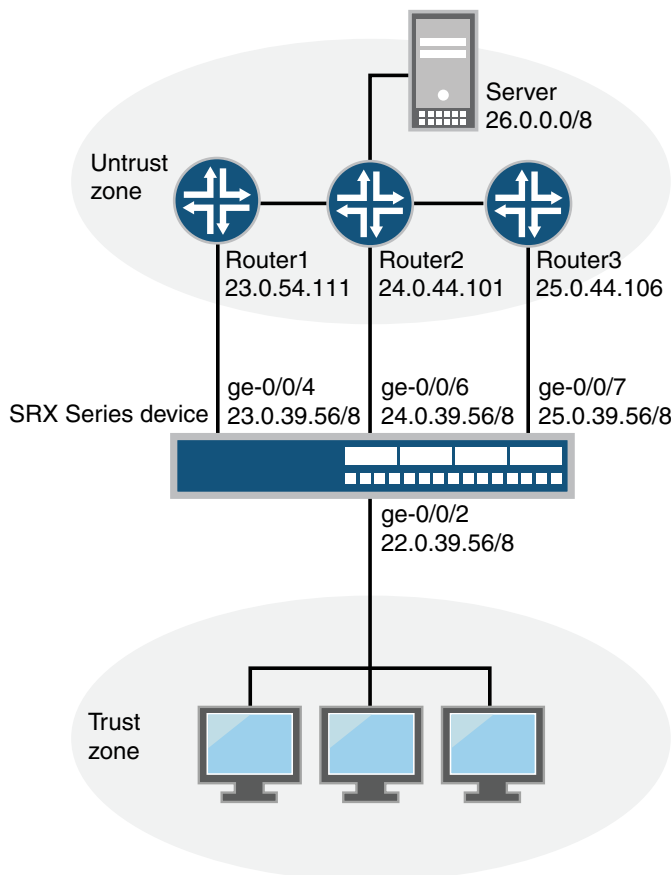
This example configures three static ECMP routes on an SRX Series device. Each static route uses a different router to reach the destination server. The interfaces to the routers are assigned to the untrust security zone. This example creates a load-balancing routing policy named **load-balancing-policy** and applies the policy to all routes exported from the routing table to the forwarding table.

### Topology

---

[Figure 15 on page 128](#) shows the topology used in this example.

Figure 15: ECMP Routes



## Configuration

### CLI Quick Configuration

To quickly configure ECMP flow-based forwarding, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```
set interfaces ge-0/0/2 unit 0 family inet address 22.0.39.56/8
set interfaces ge-0/0/4 unit 0 family inet address 23.0.39.56/8
set interfaces ge-0/0/6 unit 0 family inet address 24.0.39.56/8
set interfaces ge-0/0/7 unit 0 family inet address 25.0.39.56/8
set routing-options static route 26.0.0.0/8 next-hop 23.0.54.111
set routing-options static route 26.0.0.0/8 next-hop 24.0.44.101
set routing-options static route 26.0.0.0/8 next-hop 25.0.44.106
set security zones security-zone trust interfaces ge-0/0/2
set security zones security-zone untrust interfaces ge-0/0/4 ge-0/0/6 ge-0/0/7
set security policies from-zone trust to-zone untrust policy permit-mail match
  source-address 22.0.39.56/8
set security policies from-zone trust to-zone untrust policy permit-mail match
  destination-address 26.0.0.0/8
set security policies from-zone trust to-zone untrust policy permit-mail match application
  junos-mail
set security policies from-zone trust to-zone untrust policy permit-mail then permit
```

```
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set routing-options forwarding-table export load-balancing-policy
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the CLI User Guide.

To configure ECMP flow-based forwarding:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/2 unit 0 family inet address 22.0.39.56/8
user@host# set ge-0/0/4 unit 0 family inet address 23.0.39.56/8
user@host# set ge-0/0/6 unit 0 family inet address 24.0.39.56/8
user@host# set ge-0/0/7 unit 0 family inet address 25.0.39.56/8
```

2. Create security zones.

```
[edit security]
user@host# set zones security-zone trust interfaces ge-0/0/2
user@host# set zones security-zone untrust interfaces ge-0/0/4
user@host# set zones security-zone untrust interfaces ge-0/0/6
user@host# set zones security-zone untrust interfaces ge-0/0/7
```

3. Configure a security policy.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy permit-mail match source-address 22.0.39.56/8
user@host# set policy permit-mail match destination-address 26.0.0.0/8
user@host# set policy permit-mail match application junos-mail
user@host# set policy permit-mail then permit
```

4. Configure ECMP routes.

```
[edit routing-options]
user@host# set static route 26.0.0.0/8 next-hop 23.0.54.111
user@host# set static route 26.0.0.0/8 next-hop 24.0.44.101
user@host# set static route 26.0.0.0/8 next-hop 25.0.44.106
```

5. Create a load-balancing routing policy.

```
[edit policy-options]
user@host# set policy-statement load-balancing-policy then load-balance per-packet
```

6. Apply the routing policy to all routes exported from the routing table to the forwarding table.

```
[edit routing-options]
user@host# set forwarding-table export load-balancing-policy
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show security**, **show policy-options** and **show routing-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/2 {
```

```
    unit 0 {
      family inet {
        address 22.0.39.56/8;
      }
    }
  }
ge-0/0/4 {
  unit 0 {
    family inet {
      address 23.0.39.56/8;
    }
  }
}
ge-0/0/6 {
  unit 0 {
    family inet {
      address 24.0.39.56/8;
    }
  }
}
ge-0/0/7 {
  unit 0 {
    family inet {
      address 25.0.39.56/8;
    }
  }
}
}
user@host# show security
policies {
  from-zone trust to-zone untrust {
    policy permit-mail {
      match {
        source-address 22.0.39.56/8;
        destination-address 26.0.0.0/8;
        application junos-mail;
      }
      then {
        permit;
      }
    }
  }
}
zones {
  security-zone trust {
    interfaces {
      ge-0/0/2.0;
    }
  }
  security-zone untrust {
    interfaces {
      ge-0/0/4.0;
      ge-0/0/6.0;
      ge-0/0/7.0;
    }
  }
}
```



```

user@host# show policy-options
policy-statement load-balancing-policy {
  then {
    load-balance per-packet;
  }
}

[edit]
user@host# show routing-options
forwarding-table {
  export load-balancing-policy;
}
static {
  route 0.0.0.0/0 next-hop 10.100.37.1;
  route 26.0.0.0/8 next-hop [ 23.0.54.111 25.0.44.106 24.0.44.101 ];
}

```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

### Verifying Forwarding Table

**Purpose** Verify that the route information for all ECMP routes appears in the forwarding table.

**Action** From operational mode, enter the **show route forwarding-table** command.

```

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
...

26.0.0.0/8       user  0 23.0.54.111  rslv  0  1 ge-0/0/4.0
26.0.0.0/8       user  0 24.0.44.101  rslv  0  1 ge-0/0/6.0
26.0.0.0/8       user  0 25.0.44.106  rslv  0  1 ge-0/0/7.0
...

```

**Related Documentation**

- [Understanding ECMP Flow-Based Forwarding on page 126](#)



## CHAPTER 9

# Extended Match Conditions Configuration

- [Configuring AS Path Regular Expressions to Use as Routing Policy Match Conditions on page 133](#)
- [Defining BGP Communities and Extended Communities for Use in Routing Policy Match Conditions on page 140](#)
- [Including BGP Communities and Extended Communities in Routing Policy Match Conditions on page 146](#)
- [Using Routing Policies to Prevent Advertisement of BGP Communities to Neighbors on page 147](#)
- [Examples: Configuring BGP Communities as Routing Policy Match Conditions on page 147](#)
- [Configuring Prefix Lists for Use in Routing Policy Match Conditions on page 151](#)
- [Example: Configuring a Prefix List on page 155](#)
- [Configuring Route Lists for Use in Routing Policy Match Conditions on page 155](#)
- [Configuring Subroutines in Routing Policy Match Conditions on page 170](#)
- [Configuring Routing Policy Match Conditions Based on Routing Table Entries on page 174](#)

## Configuring AS Path Regular Expressions to Use as Routing Policy Match Conditions

A 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:

- [Configuring AS Path Regular Expressions on page 134](#)
- [How AS Path Regular Expressions Are Evaluated on page 139](#)
- [Examples: Configuring AS Path Regular Expressions on page 139](#)

## Configuring 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 [“Configuring Match Conditions in Routing Policy Terms” on page 51](#)). 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 include both of the `as-path` and `as-path-group` statements 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 65,536 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.

In Junos OS Release 9.1 and later, you can specify 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*, as well as the 2-byte AS numbers that are supported in earlier releases of the Junos OS. You can configure a value in the range from 1 through 4,294,967,295.

- **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 12 on page 136](#) 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 13 on page 136](#) shows examples of how to define regular expressions to match AS paths.

Table 12: AS Path Regular Expression Operators

Operator	Match Definition
<b>{<i>m</i>,<i>n</i>}</b>	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> .
<b>{<i>m</i>}</b>	Exactly <i>m</i> repetitions of <i>term</i> . <i>m</i> must be a positive integer.
<b>{<i>m</i>,}</b>	<i>m</i> or more repetitions of <i>term</i> . <i>m</i> must be a positive integer.
<b>*</b>	Zero or more repetitions of <i>term</i> . This is equivalent to {0,}.
<b>+</b>	One or more repetitions of <i>term</i> . This is equivalent to {1,}.
<b>?</b>	Zero or one repetition of <i>term</i> . This is equivalent to {0,1}.
<b> </b>	One of two terms on either side of the pipe.
<b>–</b>	Between a starting and ending range, inclusive.
<b>^</b>	A character at the beginning of a community attribute regular expression. This character is added implicitly; therefore, the use of it is optional.
<b>\$</b>	A character at the end of a community attribute regular expression. This character is added implicitly; therefore, the use of it is optional.
<b>( )</b>	A group of terms that are enclosed in the parentheses. Intervening space between the parentheses and the terms is ignored. If a set of parentheses is enclosed in quotation marks with no intervening space "()", it indicates a null path.
<b>[ ]</b>	Set of AS numbers. One AS number from the set must match. To specify the start and end of a range, use a hyphen (-). A caret (^) may be used to indicate that it does not match a particular AS number in the set, for example [^123].

Table 13: Examples of 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

Table 13: Examples of AS Path Regular Expressions (*continued*)

AS Path to Match	Regular Expression	Sample Matches
One through four occurrences of AS number 1234	1234{1,4}	1234 1234 1234 1234 1234 1234 1234 1234 1234 1234
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	123 124 125
	[123–125]*	Null AS path 123 124 124 125 125 125 123 124 125 123
Path whose second AS number must be 56 or 78	(. 56)   (. 78) or . (56   78)	1234 56 1234 78 9876 56 3857 78
Path whose second AS number might be 56 or 78	. (56   78)?	1234 56 52 34 56 1234 1234 78 39 794 78 2
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,}	12341234567812345678

Table 13: Examples of AS Path Regular Expressions (*continued*)

AS Path to Match	Regular Expression	Sample Matches
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 AS 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 4 9 4 5 6
AS path 5, 12, or 18	5   12   18	5 12 18

### Configuring a 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:

```
"()"
```

In the following example, locally administered AS 2 is connected to AS 1 (10.2.2.6) and AS 3. AS 3 advertises its routes to AS 2, but the administrator for AS 2 does not want to advertise AS 3 routes to AS 1 and thereby allow transit traffic from AS 1 to AS 3 through AS 2. To prevent transit traffic, the export policy **only-my-routes** is applied to AS 1. It permits advertisement of routes from AS 2 to AS 1 but prevents advertisement of routes for AS 3 (or routes for any other connected AS) 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;
  }
}
```

---

## Defining BGP Communities and Extended Communities for Use in Routing Policy Match Conditions

---

BGP communities can help you control routing policy. An example of a good use for BGP communities is unequal load balancing. When an autonomous system border router (ASBR) receives routes from directly connected external (EBGP) neighbors, the ASBR then advertises those routes to internal neighbors, using IBGP advertisements. In the IBGP advertisements, you can attach the link-bandwidth community to communicate the bandwidth of the advertised external link. This is useful when multiple external links are available, and you want to do unequal load balancing over the links. You configure the link-bandwidth extended community on all ingress links of the AS.



**NOTE:** The bandwidth information in the link-bandwidth extended community is based on the configured bandwidth of the EBGP link. It is not based on the amount of traffic on the link.

---

The Junos OS supports BGP link-bandwidth and multipath load balancing, as described in Internet draft draft-ietf-idr-link-bandwidth-01.txt, *BGP Link Bandwidth Extended Community* (expires August 2010).

To use a BGP community or extended community as a routing policy match condition, you define the community as described in the following sections:

- [Defining BGP Communities for Use in Routing Policy Match Conditions on page 141](#)
- [Defining BGP Extended Communities for Use in Routing Policy Match Conditions on page 144](#)
- [Inverting Community Matches on page 146](#)

## Defining BGP Communities for Use in Routing Policy Match Conditions

To create a named BGP 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]**

**name** identifies the community. 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** identifies one or more members of the community. Each community ID 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. You can use the following notation in specifying the AS number:
  - String of digits.
  - Asterisk (\*)—A wildcard character that matches all AS numbers. (In the definition of the community attribute, the asterisk also functions as described in [Table 14 on page 143](#).)
  - 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 regular expressions, see ["Using UNIX Regular Expressions in Community Names" on page 142](#).
- **community-value**—Identifier of the community member. It can be a number from 0 through 65,535. You can use the following notation in specifying the community ID:

- String of digits.
- Asterisk (\*)—A wildcard character that matches all community values. (In the definition of the community attribute, the asterisk also functions as described in [Table 14 on page 143](#).)
- 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 can also include one of the following well-known community names (defined in RFC 1997, *BGP Communities Attribute*) in the **community-ids** option for the **members** statement:

- **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.

---

### Using UNIX Regular Expressions in Community Names

When specifying the members of a named BGP community (in the **members [ community-ids ]** statement), 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 14 on page 143](#) lists the regular expression operators supported in community IDs. 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 15 on page 144](#) 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:** In Junos OS Release 9.1 and later, you can specify 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*, as well as the 2-byte AS numbers that are supported in earlier releases of the Junos OS.

**Table 14: 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.  If you omit the $^$ character, it is implicitly added.  We recommend explicit use of this operator for the clearest interpretation of your configuration.
$\$$	Character at the end of a community attribute regular expression.  If you omit the $\$$ character, it is implicitly added.  We recommend explicit use of this operator for the clearest interpretation of your configuration.
$[ ]$	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.
$" "$	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 15: Examples of Community Attribute Regular Expressions

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

## Defining BGP Extended Communities for Use in Routing Policy Match Conditions

To create a named BGP 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]

**name** identifies the community. 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** identifies one or more members of the community. Each community ID consists of three components, which you specify in the following format:

```
type:administrator:assigned-number
```

**type** is the type of extended community and can be either the 16-bit numerical identifier of a specific BGP extended community or one of these types:

- **bandwidth**—Sets up the bandwidth extended community. 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.

- **domain-id**—Identifies the OSPF domain from which the route originated.
- **origin**—Identifies where the route originated.
- **rt-import**—Identifies the route to install in the routing table.



**NOTE:** You must identify the route by an IP address, not an AS number.

- **src-as**—Identifies the AS from which the route originated. You must specify an AS number, not an IP address.



**NOTE:** You must identify the AS by an AS number, not an IP address.

- **target**—Identifies the destination to which the route is going.



**NOTE:** For an import policy for a VPN routing and forwarding (VRF) instance, you must include at least one route target. Additionally, you cannot use wildcard characters or regular expressions in the route target for a VRF import policy. Each value you configure for a route target for a VRF import policy must be a single value.

**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.

In Junos OS Release 9.1 and later, you can specify 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*, as well as the 2-byte AS numbers that are supported in earlier releases of the Junos OS. In plain-number format, you can configure a value in the range from 1 through 4,294,967,295. To configure a **target** or **origin** extended community that includes a 4-byte AS number in the plain-number format, append the letter “L” to the end of number. For example, a target community with the 4-byte AS number 334,324 and an assigned number of 132 is represented as **target:334324L:132**.

In Junos OS Release 9.2 and later, you can also use AS-dot notation when defining a 4-byte AS number for the **target** and **origin** extended communities. Specify two integers

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 represented in plain-number format as 65546 is represented in AS-dot notation as 1.10.

### Examples: Defining BGP Extended Communities

---

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 **10.1.1.1** and an assigned number of **20**:

```
[edit policy-options]
community test-a members [ target:10.1.1.1:20 ];
```

Configure an origin community with an administrative field of **10.1.1.1** and an assigned number of **20**:

```
[edit policy-options]
community test-a members [ origin:10.1.1.1:20 ];
```

Configure a target community with a 4-byte AS number in the administrative field of **100000** and an assigned number of **130**:

```
[edit policy-options]
community test-b members [ target:100000L:130 ];
```

#### Related Documentation

- [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)
- Configuring 4-Byte AS Numbers and BGP Extended Community Attributes in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)

## 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*]**

## Including BGP Communities and Extended Communities in Routing Policy Match Conditions

---

To include a BGP community or extended community in a routing policy match condition, include the **community** condition in the **from** statement of a policy term:

```
from {
```



```
community [ names ];
}
```

You can include this statement at the following hierarchy levels:

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

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.



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

## Using Routing Policies to Prevent Advertisement of BGP 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 BGP Communities as Routing Policy Match Conditions

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 lcl20x-high members "^20 [ 1-5 ] : 120$";
policy-statement in-customer {
  term term1 {
    from {
      protocol bgp;
      community lcl20x-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 Prefix Lists for Use in Routing Policy Match Conditions

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.

Suppose, for example, that you configure the following prefix list:

```

prefix-list bgp179 {
  apply-path "protocols bgp group <*> neighbor <*>";
}

```

This works well when all neighbors on the device are in the same address family.

When the neighbors are in different address families, for example when both IPv4 and IPv6 neighbors are configured, you can use a prefix list as follows:

```
prefix-list IPV4-BGP-NEIGHBORS {
  apply-path "protocols bgp group <*> neighbor <*:***>";
}
prefix-list IPV6-BGP-NEIGHBORS {
  apply-path "protocols bgp group <*> neighbor <*:*:*>";
}
```

One prefix list matches IPv4 addresses. The other matches IPv6 addresses. You can run the **show configuration policy-options prefix-list prefix-list name | display inheritance** command to verify the configuration.

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 16 on page 152](#).

**Table 16: Prefix List and Route List Differences**

Feature	Prefix List	Route Lists
Action	Can specify action in a <b>then</b> 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 <b>route-filter</b> match condition in a <b>from</b> statement, or to all prefixes in the list using a <b>then</b> statement.

For information about configuring route lists, see “[Configuring Route Lists for Use in Routing Policy Match Conditions](#)” on page 155.

This section includes the following information:

- [Configuring Prefix Lists on page 152](#)
- [How Prefix Lists Are Evaluated in Routing Policy Match Conditions on page 154](#)
- [Configuring Prefix List Filters on page 154](#)

## Configuring Prefix Lists

You can create a named prefix list and include it in a routing policy with the **prefix-list** match condition (described in “[Configuring Match Conditions in Routing Policy Terms](#)” on page 51).

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 Guidelines for Configuring Standard Firewall Filters and Guidelines for Applying Standard Firewall Filters.

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:** You cannot add a path element, including wildcards, after a leaf statement in the **apply-path** statement. Path elements, including wildcards, can only be used after a container statement.



**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 155](#).

## How Prefix Lists Are Evaluated in Routing Policy Match Conditions

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 for Use in Routing Policy Match Conditions*.

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 Prefix List Filters

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 17 on page 154](#).

The ***actions*** option is the action to take if the prefix list matches. It can be one or more of the actions listed in [“Configuring Flow Control Actions” on page 60](#) and [“Configuring Actions That Manipulate Route Characteristics” on page 60](#).

**Table 17: Route List Match Types for a Prefix List Filter**

Match Type	Match Condition
<b>exact</b>	The route shares the same most-significant bits (described by <b><i>prefix-length</i></b> ), and <b><i>prefix-length</i></b> is equal to the route’s prefix length.



Table 17: Route List Match Types for a Prefix List Filter (*continued*)

Match Type	Match Condition
<b>longer</b>	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.
<b>orlonger</b>	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 for Use in Routing Policy Match Conditions

A *route list* is a collection of match prefixes. When specifying a match 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 Route Lists Are Evaluated in Routing Policy Match Conditions](#)” on page 160 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:

- [Configuring Route Lists on page 156](#)
- [How Route Lists Are Evaluated in Routing Policy Match Conditions on page 160](#)
- [Route List Examples on page 163](#)

## Configuring Route Lists

To configure a route list, include one or more **route-filter** or **source-address-filter** statements:

```
route-filter destination-prefix match-type {
    actions;
}
source-address-filter source-prefix match-type {
    actions;
}
```

You can include these statements at the following hierarchy levels:

- [edit policy-options policy-statement *policy-name* term *term-name* from]
- [edit logical-systems *logical-system-name* policy-options policy-statement *policy-name* term *term-name* from]

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

The **destination-prefix** address is the IP version 4 (IPv4) or IP version 6 (IPv6) address 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.

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

**source-prefix** address is the IPv4 or IPv6 address 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 source or destination prefix. It can be one of the match types listed in [Table 18 on page 158](#). For examples of the match types and the results when presented with various routes, see [Table 19 on page 159](#).

**actions** are the actions to take if a route address matches the criteria specified for a destination match prefix (specified as part of a **route-filter** option) or for a source match prefix (specified as part of a **destination-address-filter** option). The actions can consist of one or more of the actions described in the following sections of “[Configuring Actions in Routing Policy Terms](#)” on page 58:

- [Configuring Flow Control Actions on page 60](#)
- [Configuring Actions That Manipulate Route Characteristics on page 60](#)

In a route list 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 but no actions are specified for 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 “[Route List Examples](#)” on page 163.

Table 18: Route List Match Types for a Prefix List

Match Type	Match Criteria
<b>address-mask</b> <b>netmask-value</b>	<p>All of the following are true:</p> <ul style="list-style-type: none"> <li>The bit-wise logical AND of the <b>netmask-value</b> pattern and the incoming IPv4 or IPv6 route address and the bit-wise logical AND of the <b>netmask-value</b> pattern and the <b>destination-prefix</b> address are the same. The bits set in the <b>netmask-value</b> pattern do not need to be contiguous.</li> <li>The <b>prefix-length</b> component of the incoming IPv4 or IPv6 route address and the <b>prefix-length</b> component of the <b>destination-prefix</b> address are the same.</li> </ul> <p><b>NOTE:</b> The <b>address-mask</b> routing policy match type is valid only for matching an incoming IPv4 (<b>family inet</b>) or IPv6 (<b>family inet6</b>) route address to a list of destination match prefixes specified in a <b>route-filter</b> statement.</p> <p>The <b>address-mask</b> routing policy match type enables you to match an incoming IPv4 or IPv6 route address on a configured netmask address in addition to the length of a configured destination match prefix. The length of the route address must match exactly with the length of the configured destination match prefix, as the <b>address-mask</b> match type does not support prefix length variations for a range of prefix lengths.</p> <p>When the longest-match lookup is performed on a route list, the lookup evaluates an <b>address-mask</b> match type differently from other routing policy match types. The lookup does not consider the length of the destination match prefix. Instead, the lookup considers the number of contiguous high-order bits set in the netmask value.</p> <p>For more information about this route list match type, see <a href="#">“How an Address Mask Match Type Is Evaluated” on page 161</a>.</p> <p>For example configurations showing route lists that contain the <b>address-mask</b> match type, see the following topics:</p> <ul style="list-style-type: none"> <li><a href="#">“Example: Accepting Incoming IPv4 Routes by Applying an Address Mask to the Route Address and the Destination Match Prefix” on page 167</a>.</li> <li><a href="#">“Example: Accepting Incoming IPv4 Routes with Similar Patterns But Different Prefix Lengths” on page 168</a>.</li> <li><a href="#">“Example: Evaluation of an Address Mask Match Type with Longest-Match Lookup” on page 169</a>.</li> </ul>
<b>exact</b>	<p>All of the following are true:</p> <ul style="list-style-type: none"> <li>The route address shares the same most-significant bits as the match prefix (<b>destination-prefix</b> or <b>source-prefix</b>). The number of significant bits is described by the <b>prefix-length</b> component of the match prefix.</li> <li>The <b>prefix-length</b> component of the match prefix is equal to the route's prefix length.</li> </ul>
<b>longer</b>	<p>All of the following are true:</p> <ul style="list-style-type: none"> <li>The route address shares the same most-significant bits as the match prefix (<b>destination-prefix</b> or <b>source-prefix</b>). The number of significant bits is described by the <b>prefix-length</b> component of the match prefix.</li> <li>The route's prefix length is greater than the <b>prefix-length</b> component of the match prefix.</li> </ul>

Table 18: Route List Match Types for a Prefix List (*continued*)

Match Type	Match Criteria
<b>orlonger</b>	<p>All of the following are true:</p> <ul style="list-style-type: none"> <li>The route address shares the same most-significant bits as the match prefix (<i>destination-prefix</i> or the <i>source-prefix</i>). The number of significant bits is described by the <i>prefix-length</i> component of the match prefix.</li> <li>The route's prefix length is equal to or greater than the <i>prefix-length</i> component of the configured match prefix.</li> </ul>
<b>prefix-length-range</b> <i>prefix-length2-prefix-length3</i>	<p>All of the following are true:</p> <ul style="list-style-type: none"> <li>The route address shares the same most-significant bits as the match prefix (<i>destination-prefix</i> or <i>source-prefix</i>). The number of significant bits is described by the <i>prefix-length</i> component of the match prefix.</li> <li>The route's prefix length falls between <i>prefix-length2</i> and <i>prefix-length3</i>, inclusive.</li> </ul>
<b>through</b> { <i>destination-prefix2</i>   <i>source-prefix2</i> }	<p>All of the following are true:</p> <ul style="list-style-type: none"> <li>The route address shares the same most-significant bits as the first match prefix (<i>destination-prefix</i> or <i>source-prefix</i>). The number of significant bits is described by the <i>prefix-length</i> component of the first match prefix.</li> <li>The route address shares the same most-significant bits as the second match prefix (<i>destination-prefix2</i> or <i>source-prefix2</i>). The number of significant bits is described by the <i>prefix-length</i> component of the second match prefix.</li> <li>The route's prefix length is less than or equal to the <i>prefix-length</i> component of the second match prefix.</li> </ul> <p>You do not use the <b>through</b> match type in most routing policy configurations. For an example, see <a href="#">"Example: Rejecting Routes from Specific Hosts" on page 164</a>.</p>
<b>upto prefix-length2</b>	<p>All of the following are true:</p> <ul style="list-style-type: none"> <li>The route address shares the same most-significant bits as the match prefix (<i>destination-prefix</i> or <i>source-prefix</i>). The number of significant bits is described by the <i>prefix-length</i> component of the match prefix.</li> <li>The route's prefix length falls between the <i>prefix-length</i> component of the first match prefix and <i>prefix-length2</i>.</li> </ul>

Table 19: Match Type Examples

Prefix	192.168.0.0/16 exact	192.168.0.0/16 longer	192.168.0.0/16 orlonger	192.168.0.0/16 upto /24	192.168.0.0/16 through 192.168.0.0/20	192.168.0.0/16 through 192.168.0.0/20	192.168.0.0/16 through 192.168.0.0/25
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	Match

Table 19: Match Type Examples (*continued*)

Prefix	192.168.16.0/24 exact	192.168.16.0/24 longer	192.168.16.0/24 longer	192.168.16.0/24 upto /24	192.168.16.0/24 longer	192.168.16.0/24 longer	192.168.16.0/24 longer
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	—	—
192.168.22.40/19	—	Match	Match	Match	Match	—	Match
10.169.1.0/24	—	—	—	—	—	—	—
10.170.0.0/16	—	—	—	—	—	—	—

## How Route Lists Are Evaluated in Routing Policy Match Conditions

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* components of the configured match prefix only, and not the *match-type* component. 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 for the candidate route **192.168.1.0/24** 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 actions 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.



**NOTE:** When a route source address is evaluated against a match criteria that uses the **address-mask** match type, both steps of the evaluation include the configured netmask value. For more information, see [“How an Address Mask Match Type Is Evaluated” on page 161](#).

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 18](#).



**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.

### How an Address Mask Match Type Is Evaluated

The **address-mask** routing policy match type enables you to match incoming IPv4 or IPv6 route addresses on a configured netmask value in addition to the length of a configured destination match prefix. During route list evaluation, an **address-mask** match type is processed differently from other routing policy match types, taking into consideration the configured netmask value:

1. When a longest-match lookup evaluates an **address-mask** routing policy match type, the **prefix-length** component of the configured match prefix is not considered. Instead,

the lookup considers the number of contiguous high-order bits set in the configured netmask value.

2. When an incoming IPv4 or IPv6 route address is evaluated against a route filter match criteria that uses the **address-mask** routing policy match type, the match succeeds if the following values are identical:
  - The bit-wise logical AND of the configured netmask value and the incoming IPv4 or IPv6 route address
  - The bit-wise logical AND of the configured netmask value and the configured destination match prefix

For an example configuration of a route list that contains two **address-mask** match types, see [“Example: Evaluation of an Address Mask Match Type with Longest-Match Lookup” on page 169](#).

---

### 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 18](#).) 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 or longer** from the route list in this term and move it to another term where it is the only prefix or the longest prefix in the list.

## Route List Examples

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 18](#).

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

- [Example: Rejecting Routes with Specific Destination Prefixes and Mask Lengths on page 163](#)
- [Example: Rejecting Routes with a Mask Length Greater than Eight on page 164](#)
- [Example: Rejecting Routes with Mask Length Between 26 and 29 on page 164](#)
- [Example: Rejecting Routes from Specific Hosts on page 164](#)
- [Example: Accepting Routes with a Defined Set of Prefixes on page 165](#)
- [Example: Rejecting Routes with a Defined Set of Prefixes on page 165](#)
- [Example: Rejecting Routes with Prefixes Longer than 24 Bits on page 166](#)
- [Example: Rejecting PIM Multicast Traffic Joins on page 166](#)
- [Example: Rejecting PIM Traffic on page 166](#)
- [Example: Accepting Incoming IPv4 Routes by Applying an Address Mask to the Route Address and the Destination Match Prefix on page 167](#)
- [Example: Accepting Incoming IPv4 Routes with Similar Patterns But Different Prefix Lengths on page 168](#)
- [Example: Evaluation of an Address Mask Match Type with Longest-Match Lookup on page 169](#)

### 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;
      }
      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 Multicast Protocols Configuration Guide.

#### Example: Accepting Incoming IPv4 Routes by Applying an Address Mask to the Route Address and the Destination Match Prefix

Accept incoming IPv4 routes with a destination prefix of **10.1.0/24** and the third byte an even number from 0 to 14, inclusive:

```
[edit]
policy-options {
  policy-statement from_customer_a {
    term term_1 {
      from {
        route-filter 10.1.0.0/24 address-mask 255.255.241.0;
      }
      then {
        ...
        reject;
      }
    }
  }
}
```

The route list in routing policy term **term\_1** matches the following incoming IPv4 route addresses:

- 10.1.0.0/24
- 10.1.2.0/24
- 10.1.4.0/24
- 10.1.6.0/24
- 10.1.8.0/24
- 10.1.10.0/24
- 10.1.12.0/24
- 10.1.14.0/24

The bit-wise logical AND of the netmask value and the candidate route address must match the bit-wise logical AND of the netmask value and the match prefix address. That is, where the netmask bit pattern **255.255.241.0** contains a set bit, the incoming IPv4 route address being evaluated must match the value of the corresponding bit in the destination prefix address **10.1.0.0/24**.

- The first two bytes of the netmask value are binary **1111 1111 1111 1111**, which means that a candidate route address will fail the match if the first two bytes are not **10.1**.
- The third byte of the netmask value is binary **1111 0001**, which means that a candidate route address will fail the match if the third byte is greater than 15 (decimal), an odd number, or both.
- The prefix length of the match prefix address is **24** (decimal), which means that a candidate route address will fail the match if its prefix length is not exactly **24**.

As an example, suppose that the candidate route address being tested in the policy is **10.1.8.0/24** (binary **0000 1010 0000 0001 0000 1000**).

1. When the netmask value is applied to this candidate route address, the result is binary **0000 1010 0000 0001 0000 0000**.
2. When the netmask value is applied to the configured destination prefix address, the result is also binary **0000 1010 0000 0001 0000 0000**.
3. Because the results of both AND operations are the same, the match continues to the second match criteria.
4. Because the prefix lengths of the candidate address and the configured destination prefix address are the same (24 bits), the match succeeds.

As another example, suppose that the candidate route address being tested in the policy is **10.1.3.0/24** (binary **0000 1010 0000 0001 0000 0011**).

1. When the netmask value is applied to this candidate route address, the result is binary **0000 1010 0000 0001 0000 0001**.
2. However, when the netmask value is applied to the configured destination prefix address, the result is binary **0000 1010 0000 0001 0000 0000**.
3. Because the results of the two AND operations are different (in the third byte), the match fails.

### Example: Accepting Incoming IPv4 Routes with Similar Patterns But Different Prefix Lengths

---

Accept incoming IPv4 route addresses of the form **10.\*1/24** or **10.\*1./32**:

```
[edit]
policy-options {
  policy-statement from_customer_b {
    term term_2 {
      from {
        route-filter 10.0.1.0/24 address-mask 255.0.255.0;
        route-filter 10.0.1.0/32 address-mask 255.0.255.0;
```

```

    }
    then {
        ...
        reject;
    }
}
}
}

```

The route filter match criteria **10.0.1.0/24 address-mask 255.0.255.0** matches an incoming IPv4 route address of the form **10.\*1/24**. The route's prefix length must be exactly 24 bits long, and any value is acceptable in the second byte.

The route filter match criteria **10.0.1.0/32 address-mask 255.0.255.0** matches an incoming IPv4 route address of the form **10.\*1.\*32**. The route's prefix length must be exactly 32 bits long, and any value is acceptable in the second byte and the fourth byte.

### Example: Evaluation of an Address Mask Match Type with Longest-Match Lookup

This example illustrates how a longest-match lookup evaluates a route list that contains two **address-mask** match types. Consider the route list configured in the routing policy term **term\_3** below:

```

[edit]
policy-options {
  policy-statement from_customer_c {
    term term_3 {
      from {
        route-filter 10.0.1.0/24 address-mask 255.0.255.0;
        route-filter 10.0.2.0/24 address-mask 255.240.255.0;
      }
      then {
        ...
      }
    }
  }
}

```

Suppose that the incoming IPv4 route source address **10.1.1.0/24** is tested against the route list configured in the policy term **term\_3**:

1. The longest-match lookup tree for routing policy term **term\_3** contains two match prefixes: one prefix for **10.0.1.0/24 address-mask 255.0.255.0** and one prefix for **10.0.2.0/24 address-mask 255.240.255.0**. When searching the tree for the longest-prefix match for a candidate, the longest-match lookup considers the number of contiguous high-order bits in the configured **netmask-value** instead of the length of the configured **destination-prefix**:
  - For the first route filter match criteria, the longest-match lookup entry is **10.0.0.0/8** because the netmask value contains 8 contiguous high-order bits.
  - For second route filter match criteria, the longest-match lookup entry is **10.0.0.0/12** because the netmask value contains 12 contiguous high-order bits.

For the candidate route address **10.1.1.0/24**, the longest-match lookup returns the tree entry **10.0.0.0/12**, which corresponds to the route filter match criteria **10.0.2.0/24 address-mask 255.240.255.0**.

2. Now that the longest-match prefix in **term\_3** has been identified for the candidate route address, the candidate route address is evaluated against the route filter match criteria **10.0.2.0/24 address-mask 255.240.255.0**:
  - a. To test the incoming IPv4 route address **10.1.1.0/24**, the netmask value **255.240.255.0** is applied to **10.1.1.0/24**. The result is **10.0.1.0**.
  - b. To test the configured destination prefix address **10.0.2.0/24**, the netmask value **255.240.255.0** is applied to **10.0.2.0/24**. The result is **10.0.2.0**.
  - c. Because the results are different, the route filter match fails. No actions, whether specified with the match criteria or with the **then** statement, are taken. The incoming IPv4 route address is not evaluated against any other match criteria.

## Configuring Subroutines in Routing Policy Match Conditions

---

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

For configuration instructions and examples, see the following section:

- [Configuring Subroutines on page 170](#)
- [Example: Configuring a Subroutine on page 173](#)

### Configuring 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 “[Configuring Match Conditions in Routing Policy Terms](#)” on page 51) in the **from** or **to** statement of another routing policy:

```
[edit]
policy-options {
  policy-statement subroutine-policy-name {
    term term-name {
      from {
        match-conditions;
        prefix-list name;
        route-filter destination-prefix match-type <actions>;
        source-address-filter source-prefix match-type <actions>;
      }
      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. The result is that no prefixes 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 28](#).

### Possible Consequences of Termination Actions in Subroutines

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 18](#).)

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.
- 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 Routing Policy Match Conditions Based on Routing Table Entries

In addition to defining match conditions in the **from** statement of a policy, you can use the **condition** statement to define conditions used during 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.

**Related  
Documentation**

- **show policy conditions** in the Junos OS Operational Mode Commands



# Extended Actions Configuration

- [Prepending AS Numbers to BGP AS Paths on page 177](#)
- [Example: Configuring a Routing Policy to Prepend the AS Path on page 178](#)
- [Adding AS Numbers to BGP AS Paths on page 180](#)
- [Using Routing Policies to Damp BGP Route Flapping on page 181](#)
- [Configuring Per-Packet Load Balancing on page 186](#)
- [Configuring Load Balancing Based on MPLS Labels on DPC I-Chip-Based Hardware on page 189](#)
- [Configuring Load Balancing for Ethernet Pseudowires on page 193](#)
- [Configuring Load Balancing Based on MAC Addresses on page 194](#)
- [Configuring VPLS Load Balancing Based on IP and MPLS Information on page 195](#)
- [Configuring VPLS Load Balancing on MX Series 3D Universal Edge Routers on page 196](#)
- [Example: Configuring Multicast Load Balancing over Aggregated Ethernet Links on page 197](#)

## Prepending AS Numbers to BGP AS Paths

---

You can *prepend* one or more autonomous system (AS) numbers at the beginning of an AS path. The AS numbers are added at the beginning of the path after the actual AS number from which the route originates has been added to the path. Prepending an AS path makes a shorter AS path look longer and therefore less preferable to BGP.

In Junos OS Release 9.1 and later, you can specify 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*, as well as the 2-byte AS numbers that are supported in earlier releases of the Junos OS. In plain-number format, you can configure a value in the range from 1 through 4,294,967,295.

If you have a router that does not support 4-byte AS numbers in the AS path, the prepended AS number displayed in the AS path is the AS\_TRANS number, AS 23456. To display the route details, use the *show route* command. For more information, see [Prepending 4-Byte AS Numbers in an AS Path in the \*Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview\*](#).

In the following 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 the configuration for AS 1, 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 "1111";
    }
  }
}
```

**Related  
Documentation**

- 4-Byte Autonomous System Numbers Overview in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)
- Prepending 4-Byte AS Numbers in an AS Path in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)

---

## Example: Configuring a Routing Policy to Prepend the AS Path

This example shows how to configure a routing policy to prepend the AS path.

- [Requirements on page 178](#)
- [Overview on page 178](#)
- [Configuration on page 178](#)
- [Verification on page 180](#)

### Requirements

Before you begin, configure router interfaces and configure routing protocols, as explained in [“Routing Policies Configuration Overview” on page 10](#).

### Overview

In this example, you create a routing policy called `prependpolicy1` and a term called `prependterm1`. The routing policy prepends the AS numbers 1111 to routes that are greater than or equal to 172.16.0.0/12, 192.168.0.0/16, and 10.0.0.0/8. The policy is applied as an import policy to all BGP routes and is evaluated when routes are imported to the routing table.

### Configuration

**CLI Quick  
Configuration**

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



```

set policy-options policy-statement prependpolicy1 term prependterm1 from route-filter
  172.16.0.0/12 orlonger
set policy-options policy-statement prependpolicy1 term prependterm1 from route-filter
  192.168.0.0/16 orlongerr
set policy-options policy-statement prependpolicy1 term prependterm1 from route-filter
  10.0.0.0/8 orlonger
set policy-options policy-statement prependpolicy1 term prependterm1 then
  as-path-prepend 200
set policy-options policy-statement prependpolicy1 term prependterm1 then accept
set protocols bgp import prependpolicy1

```

### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the CLI User Guide.

To create a routing policy that prepends AS numbers to multiple routes:

1. Create the routing policy conditions.

```

[edit policy-options policy-statement prependpolicy1 term prependterm1]
user@R2# set from route-filter 172.16.0.0/12 orlonger
user@R2# set from route-filter 192.168.0.0/16 orlongerr
user@R2# set from route-filter 10.0.0.0/8 orlonger

```

2. Create the routing policy actions, including the AS number to prepend.

```

[edit policy-options policy-statement prependpolicy1]
user@R2# set then as-path-prepend 200
user@R2# set then accept

```



**NOTE:** If you enter multiple numbers, you must separate each number with a space. Enclose the numbers in double quotation marks.

3. Apply the policy as an import policy for all BGP routes in the local routing table.

Apply the policy as an export policy for all BGP routes advertised to neighbors.

```

[edit protocols bgp]
user@host# set import prependpolicy1

```

```

user@host# set export prependpolicy1

```



**NOTE:** You can refer to the same routing policy one or more times in the same or different import statement.

**Results** Confirm your configuration by entering the **show policy-options** and **show protocols bgp** commands from configuration mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show policy-options
policy-statement prependpolicy1 {
  term prependterm1 {
    from {
      route-filter 172.16.0.0/12 orlonger;
      route-filter 192.168.0.0/16 orlonger;
      route-filter 10.0.0.0/8 orlonger;
    }
    then as-path-prepend 200;
    then accept;
  }
}

user@host# show protocols bgp
export prependpolicy1;
```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying the AS Numbers Are Prepended on page 180](#)

---

### Verifying the AS Numbers Are Prepended

<b>Purpose</b>	Verify that the AS path is prepended with the specified AS number.
<b>Action</b>	From operational mode, enter the <b>show route</b> command.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Routing Policies Configuration Overview on page 10</a></li><li>• <a href="#">Prepending AS Numbers to BGP AS Paths on page 177</a></li></ul>

---

## Adding AS Numbers to BGP AS Paths

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 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.

## Using Routing Policies to Damp BGP Route Flapping

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 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 [“Configuring Actions That Manipulate Route Characteristics” on page 60](#)). For the damping routing policy to work, you also must enable BGP route flap damping.

The following sections discuss the following topics:

- [Configuring BGP Flap Damping Parameters on page 181](#)
- [Specifying BGP Flap Damping as the Action in Routing Policy Terms on page 184](#)
- [Disabling Damping for Specific Address Prefixes on page 184](#)
- [Example: Configuring BGP Flap Damping on page 185](#)

## Configuring BGP 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 20 on page 182](#).

**Table 20: Damping Parameters**

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



**NOTE:** For the half-life, configure a value that is less than the max-suppress. If you do not, the configuration is rejected.

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.



**NOTE:** For the max-suppress, configure a value that is greater than the half-life. If you do not, the configuration is rejected.

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,  $\epsilon_c$ , which is the maximum merit that a flapping route can collect, is calculated using the following formula:

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

$\epsilon_r$  is the figure-of-merit reuse threshold,  $t$  is the maximum hold-down time in minutes, and  $\lambda$  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,  $\epsilon_c$ . 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 routing device, use the **show route detail** command.

## Specifying BGP Flap Damping as the Action in Routing Policy Terms

To BGP flap damping as the action in a routing policy term, include the **damping** statement and the name of the configured damping parameters either as an option of the **route-filter** statement at the **[edit policy-options policy-statement *policy-name* term *term-name* from]** hierarchy level:

```
[edit policy-options policy-statement policy-name term term-name from]
route-filter destination-prefix match-type {
  damping damping-parameters;
}
```

or at the **[edit policy-options policy-statement *policy-name* term *term-name* then]** hierarchy level:

```
[edit policy-options policy-statement policy-name term term-name then]
damping damping-parameters;
```

## Disabling Damping for Specific Address Prefixes

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 Damping for a Specific Address Prefix

In this routing policy example, although damping is enabled for the peer, the **damping none** statement specifies that damping be disabled for prefix **10.0.0.0/8** in **Policy-A**. This route is not damped because the routing policy statement named **Policy-A** filters on the prefix **10.0.0.0/8** 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 Per-Packet Load Balancing

To configure per-packet load balancing as described in [“Overview of Per-Packet Load Balancing” on page 37](#), include the **load-balance per-packet** statement either as an option of the **route-filter** statement at the **[edit policy-options policy-statement *policy-name* term *term-name* from]** hierarchy level:

```
[edit policy-options policy-statement policy-name term term-name from]
route-filter destination-prefix match-type {
  load-balance per-packet;
```



```
}
```

or at the `[edit policy-options policy-statement policy-name term term-name then]` hierarchy level:

```
[edit policy-options policy-statement policy-name term term-name then]
load-balance per-packet;
```

To complete the configuration you must apply the routing policy to routes exported from the routing table to the forwarding table, by including the policy name in the list specified by 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.

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;
}
```

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
- IP type of service

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.

With M Series (with the exception of the M120 router) and T Series routers, the first fragment is mapped to the same load-balanced destination as the unfragmented packets. The other fragments can be mapped to other load-balanced destinations.

For the M120 router only, all fragments are mapped to the same load-balanced destination. This destination is not necessarily the same as that for unfragmented packets.

By default, or if you include 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, IP 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

## Per-Packet Load Balancing Examples

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;
  }
}
```

## Configuring Load Balancing Based on MPLS Labels on DPC I-Chip-Based Hardware

Juniper Networks routers can load-balance on a per-packet basis in MPLS. Load balancing can be performed on information in both the IP header and on up to three MPLS labels, providing a more uniform distribution of MPLS traffic to next hops. This feature is enabled on supported platforms by default and requires no configuration.



**NOTE:** MPC cards do not support the regular hash key configuration. For the MPC-based hash key configuration to be effective, you need an [enhanced-hash-key](#) configuration.

Load balancing is used to evenly distribute traffic when the following conditions apply:

- There are multiple equal-cost next hops over different interfaces to the same destination.
- There is a single next hop over an aggregated interface.

By default, when load balancing is used to help distribute traffic, Junos OS employs a hash algorithm to select a next-hop address to install into the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is reselected by means of the hash algorithm. You can configure how the hash algorithm is used to load-balance traffic across a set of equal-cost label switched paths (LSPs).

An LSP tends to load-balance its placement by randomly selecting one of the equal-cost next hops and using it exclusively. The random selection is made independently at each transit router, which compares Interior Gateway Protocol (IGP) metrics alone. No consideration is given to bandwidth or congestion levels.

To load-balance based on the MPLS label information, configure the **family mpls** statement:

```
[edit forwarding-options hash-key]
family mpls {
  all-labels;
  label-1;
  label-2;
  label-3;
  no-labels;
  no-label-1-exp;
  payload {
    ether-pseudowire;
    ip {
      disable;
      layer-3-only;
      port-data {
        destination-lsb;
        destination-msb;
        source-lsb;
        source-msb;
      }
    }
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* forwarding-options hash-key]
- [edit forwarding-options hash-key]

This feature applies to aggregated Ethernet and aggregated SONET/SDH interfaces as well as multiple equal-cost MPLS next hops. In addition, on the T Series, MX Series, M120, and M320 routers only, you can configure load balancing for IPv4 traffic over Layer 2 Ethernet pseudowires. You can also configure load balancing for Ethernet pseudowires based on IP information. The option to include IP information in the hash key provides support for Ethernet circuit cross-connect (CCC) connections.

[Table 21 on page 190](#) provides detailed information about all of the possible MPLS LSP load-balancing options.

**Table 21: MPLS LSP Load Balancing Options**

Statement	Supported Platforms	MPLS LSP Load Balancing Options
<b>all-labels</b>	PTX Series	Up to eight MPLS labels are included in the hash key to identify the uniqueness of a flow in the Packet Forwarding Engine. This value is set by default.
<b>label-1</b>	M Series, MX Series, T Series	Include the first label in the hash key. Use this option for single label packets.
<b>label-2</b>	M Series, MX Series, T Series	Include the second label in the hash key. You must also configure the <b>label-1</b> option. The entire first label and the first 16 bits of the second label are used in the hash key.

Table 21: MPLS LSP Load Balancing Options (*continued*)

Statement	Supported Platforms	MPLS LSP Load Balancing Options
<b>label-3</b>	M Series, MX Series, T Series	Include the third label in the hash key. You must also configure the <b>label-1</b> option and the <b>label-2</b> option.
<b>no-labels</b>	All	Excludes MPLS labels from the hash key.
<b>no-label-1-exp</b>	M Series, MX Series, T Series	Excludes the EXP bit of the top label from the hash key. You must also configure the <b>label-1</b> option.  For Layer 2 VPNs, the router could encounter a packet reordering problem. When a burst of traffic pushes the customer traffic bandwidth to exceed its limits, the traffic might be affected in mid flow. Packets might be reordered as a result. By excluding the EXP bit from the hash key, you can avoid this reordering problem.
<b>payload</b>	All	Allows you to configure which parts of the IP packet payload to include in the hash key. For the PTX Series Packet Transport Switch, this value is set by default.
<b>disable</b>	PTX Series	Exclude IP payload from the hash key.
<del>ether-pseudowire</del>	M120, M320, MX Series, T Series	Load-balance IPv4 traffic over Layer 2 Ethernet pseudowires.
<b>ip</b>	All	Include the IPv4 or IPv6 address in the hash key. You must also configure either <b>label-1</b> or <b>no-labels</b> .
<b>layer-3-only</b>	All	Include only the Layer 3 IP information in the hash key. Excludes all of the <b>port-data</b> bytes from the hash key.
<b>port-data</b>	M Series, MX Series, T Series	Include the source and destination port field information. By default, the most significant byte and least significant byte of the source and destination port fields are used in the hash key. To select specific bytes to use in the hash key, include one or more of the <b>source-msb</b> , <b>source-lsb</b> , <b>destination-msb</b> , and <b>destination-lsb</b> options at the <b>[edit forwarding-options hash-key family mpls payload ip port-data]</b> hierarchy level. To prevent all four bytes from being hashed, include the <b>layer-3-only</b> statement at the <b>[edit forwarding-options hash-key family mpls payload ip]</b> hierarchy level.
<del>destination-lsb</del>	M Series, MX Series, T Series	Include the least significant byte of the destination port in the hash key. Can be combined with any of the other <b>port-data</b> options.
<del>destination-msb</del>	M Series, MX Series, T Series	Include the most significant byte of the destination port in the hash key. Can be combined with any of the other <b>port-data</b> options.
<b>source-lsb</b>	M Series, MX Series, T Series	Include the least significant byte of the source port in the hash key. Can be combined with any of the other <b>port-data</b> options.
<b>source-msb</b>	M Series, MX Series, T Series	Include the most significant byte of the source port in the hash key. Can be combined with any of the other <b>port-data</b> options.

The following examples illustrate ways in which you can configure MPLS LSP load balancing:

- To include the IP address as well as the first label in the hash key:
  - For M Series, MX Series, and T Series routers, configure the **label-1** statement and the **ip** option for the **payload** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level:
- For PTX Series Packet Transport Switches, the **all-labels** and **ip payload** options are configured by default, so no configuration is necessary.
- (M320 and T Series routers only) To include the IP address as well as both the first and second labels in the hash key, configure the **label-1** and **label-2** options and the **ip** option for the **payload** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
payload {
  ip;
}
```

```
[edit forwarding-options hash-key family mpls]
label-1;
label-2;
payload {
  ip;
}
```



**NOTE:** You can include this combination of statements on M320 and T Series routers only. If you include them on an M Series Multiservice Edge Router, only the first MPLS label and the IP payload are used in the hash key.

- For T Series routers, ensure proper load balancing by including the **label-1**, **label-2**, and **label-3** options at the **[edit forwarding-options hash-key family mpls]** hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
label-2;
label-3;
```

- (M Series, MX Series, and T Series routers only) For Layer 2 VPNs, the router could encounter a packet reordering problem. When a burst of traffic pushes the customer traffic bandwidth to exceed its limits, the traffic might be affected in mid flow. Packets might be reordered as a result. By excluding the EXP bit from the hash key, you can avoid this reordering problem. 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;
}
```

Related  
Documentation

- [Configuring Load Balancing for Ethernet Pseudowires on page 193](#)

## Configuring Load Balancing for Ethernet Pseudowires

You can configure load balancing for IPv4 traffic over Layer 2 Ethernet pseudowires. You can also configure load balancing for Ethernet pseudowires based on IP information. The option to include IP information in the hash key provides support for Ethernet circuit cross-connect (CCC) connections.



**NOTE:** This feature is supported only on M120, M320, MX Series, and T Series routers.

To configure load balancing for IPv4 traffic over Layer 2 Ethernet pseudowires, include the **ether-pseudowire** statement at the **[edit forwarding-options hash-key family mpls payload]** hierarchy level:

```
[edit forwarding-options]
hash-key {
  family mpls {
    (label-1 | no-labels);
    payload {
      ether-pseudowire;
    }
  }
}
```



**NOTE:** You must also configure either the **label-1** or the **no-labels** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level.

You can also configure load balancing for Ethernet pseudowires based on IP information. This functionality provides support for load balancing for Ethernet cross-circuit connect (CCC) connections. To include IP information in the hash key, include the **ip** statement at the **[edit forwarding-options hash-key family mpls payload]** hierarchy level:

```
[edit forwarding-options]
hash-key {
  family mpls {
    (label-1 | no-labels);
    payload {
      ip;
    }
  }
}
```



**NOTE:** You must also configure either the `label-1` or `no-labels` statement at the `[edit forwarding-options hash-key family mpls]` hierarchy level.

You can configure load balancing for IPv4 traffic over Ethernet pseudowires to include only Layer 3 IP information in the hash key. To include only Layer 3 IP information, include the `layer-3-only` option at the `[edit forwarding-options family mpls hash-key payload ip]` hierarchy level:

```
[edit forwarding-options]
hash-key {
  family mpls {
    (label-1 | no-labels);
    payload {
      ip {
        layer-3-only;
      }
    }
  }
}
```



**NOTE:** You must also configure either the `label-1` or `no-labels` statement at the `[edit forwarding-options hash-key family mpls]` hierarchy level.

## Configuring Load Balancing Based on MAC Addresses

The hash key mechanism for load-balancing uses Layer 2 media access control (MAC) information such as frame source and destination address. To load-balance traffic based on Layer 2 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:** Any packets that have the same source and destination address will be sent over the same path.



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





**NOTE:** Aggregated Ethernet member links will now use the physical MAC address as the source MAC address in 802.3ah OAM packets.

**Related Documentation**

- Junos OS VPNs Configuration Guide

## Configuring VPLS Load Balancing Based on IP and MPLS Information

In Junos OS Release 9.4 and later, you can configure load balancing for VPLS traffic to have the hash key include IP information and MPLS labels on the M120 and M320 routers only. In earlier Junos OS Releases, you can configure load balancing based only on Layer 2 information. In Junos OS Release 9.5 and later, you can configure load balancing for VPLS traffic based on Layer 3 IP and Layer 4 information on MX Series routers only. For more information, see [“Configuring VPLS Load Balancing on MX Series 3D Universal Edge Routers” on page 196](#).

For IPv4 traffic, only the IP source and destination addresses are included in the hash key. For MPLS and IPv4 traffic, one or two MPLS labels and IPv4 source and destination addresses are included. For MPLS Ethernet pseudowires, only one or two MPLS labels are included in the hash key.



**NOTE:** VPLS load balancing based on MPLS labels and IP information is supported only on the M120 and M320 routers. In Junos OS Release 9.5 and later, on MX Series routers only, you can configure VPLS load balancing based on IP and Layer 4 information.

To optimize VPLS flows across multiple paths based on IP and MPLS information, include the **family multiservice** statement at the **[edit forwarding-options hash-key]** hierarchy level:

```
[edit forwarding-options hash-key]
family multiservice {
  label-1;
  label-2;
  payload {
    ip {
      layer-3-only;
    }
  }
}
```

To use the first MPLS label in the hash key, include the **label-1** statement:

```
[edit forwarding-options hash-key family multiservice]
label-1;
```

To use the second MPLS label, include both the **label-1** and **label-2** statements:

```
[edit forwarding-options hash-key family multiservice]
```

```
label-1;  
label-2;
```

To use the packet's IPv4 payload in the hash key, include the **payload** and **ip** statements:

```
[edit forwarding-options hash-key family multiservice]  
payload {  
  ip;  
}
```

---



**NOTE:** Only IPv4 is supported.

---

To include only Layer 3 information from the IPv4 payload, specify the **layer-3-only** option to the **payload ip** statement:

```
[edit forwarding-options hash-key family multiservice]  
payload {  
  ip {  
    layer-3-only;  
  }  
}
```

To use the first and second MPLS labels and the packet's IP payload in the hash key, include the **label-1**, **label-2**, and **payload ip** statements:

```
[edit forwarding-options hash-key family multiservice]  
label-1;  
label-2;  
payload {  
  ip;  
}
```

---

## Configuring VPLS Load Balancing on MX Series 3D Universal Edge Routers

---

In Junos OS Release 9.5 and later, on MX Series routers, you can configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers inside the frame payload. You can also configure VPLS load balancing based on IP and MPLS information on M120 and M320 routers only. For more information, see [“Configuring VPLS Load Balancing Based on IP and MPLS Information” on page 195](#).

You can configure load balancing on MX Series routers based on Layer 3 or Layer 4 information or both.

To configure VPLS load balancing on the MX Series router to include either Layer 3 IP information or Layer 4 headers or both:

1. Include the **payload** statement at the **[edit forwarding-options hash-key family multiservice]** hierarchy level.
2. Include the **ip** statement at the **[edit forwarding-options hash-key family multiservice payload]** hierarchy level.

To configure VPLS load balancing to include the Layer 3 information:

1. Include the **layer-3** statement at the **[edit forwarding-options hash-key family multiservice payload ip]** hierarchy level.
2. Include the **source-ip-only** statement at the **[edit forwarding-options hash-key family multiservice payload ip layer-3]** hierarchy level to include information about the IP source address only in the hash key.
3. Include **destination-ip-only** statement at the **[edit forwarding-options hash-key family multiservice payload ip layer-3]** hierarchy level to include information about the IP destination address only in the hash key.



**NOTE:** You can configure either the **source-ip-only** or the **destination-ip-only** statements at a time, not both. They are mutually exclusive.

To configure VPLS load balancing to include Layer 4 information:

- Include the **layer-4** statement at the **[edit forwarding-options hash-key family multiservice payload ip]** hierarchy level.

The following example shows load balancing configured to use the source Layer 3 IP address option and Layer 4 header fields as well as the source and destination MAC addresses:

```
[edit forwarding-options hash-key]
family multiservice {
  source-mac;
  destination-mac;
  payload {
    ip {
      layer-3 {
        source-ip-only;
      }
      layer-4;
    }
  }
}
```

- Related Documentation**
- [family multiservice on page 320](#)
  - [hash-key on page 331](#)

## Example: Configuring Multicast Load Balancing over Aggregated Ethernet Links

This example shows how to configure point-to-multipoint (P2MP) LSPs to load-balance across aggregated Ethernet links. The load balancing applies to all traffic types, including multicast. Feature parity for multicast load balancing of point-to-multipoint LSPs over

aggregated Ethernet child links on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

- [Requirements on page 198](#)
- [Overview on page 198](#)
- [Configuration on page 199](#)
- [Verification on page 209](#)

## Requirements

Before you begin:

- Configure the router interfaces. See the Junos® OS Network Interfaces.
- Configure an interior gateway protocol or static routing. See the Junos OS Routing Protocols Configuration Guide.

## Overview

This example shows a sample topology and configuration to perform the following tasks:

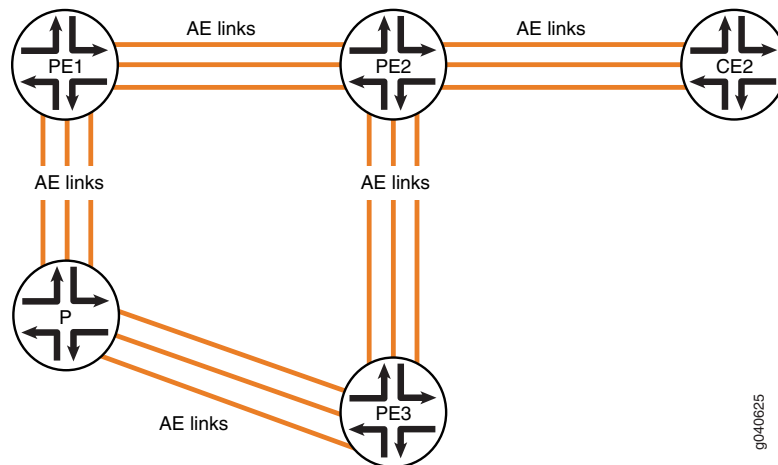
- Load balancing VPLS multicast traffic over link aggregation
- Load balancing P2MP multicast traffic over link aggregation
- Re-load balancing after a change in the nexthop topology

Nexthop topology changes might include but are not limited to:

- L2 membership change in the link aggregation
- Indirect nexthop change
- Composite nexthop change

[Figure 16 on page 198](#) shows the topology for this example. The example includes the configuration for PE1 and PE2.

**Figure 16: Multicast Load Balancing over Aggregated Ethernet Links**



9040625

## Configuration

**CLI Quick Configuration** To quickly configure load balancing across aggregated Ethernet links, copy the following commands into a text file, remove any line breaks, and then paste the commands into the CLI.

PE1:

```
[edit]
set forwarding-options hash-key family multiservice source-mac
set forwarding-options hash-key family multiservice destination-mac
set forwarding-options hash-key family multiservice payload ip layer-3
set interfaces ge-0/0/6 gigether-options 802.3ad ae0
set interfaces ge-0/1/6 gigether-options 802.3ad ae0
set interfaces ge-0/2/2 encapsulation ethernet-vpls
set interfaces ge-0/2/2 unit 0 family vpls
set interfaces ge-0/2/3 gigether-options 802.3ad ae0
set interfaces ge-0/2/6 gigether-options 802.3ad ae0
set interfaces ge-0/3/0 gigether-options 802.3ad ae0
set interfaces ge-0/3/1 gigether-options 802.3ad ae0
set interfaces ge-0/3/6 gigether-options 802.3ad ae0
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/2/6 unit 0 family inet address 13.1.1.2/30
set interfaces ae0 unit 0 family inet address 11.11.11.1/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set policy-options policy-statement exp-to-fwd term a from community grn-com
set policy-options policy-statement exp-to-fwd term a then install-nexthop lsp PE1-to-PE2
set policy-options policy-statement exp-to-fwd term a then accept
set policy-options community grn-com members target:65000:1
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path PE1-to-PE2 to 10.255.19.77
set protocols mpls label-switched-path PE1-to-PE3 to 10.255.19.79
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.71.214
set protocols bgp group int family inet any
set protocols bgp group int family l2vpn signaling
set protocols bgp group int neighbor 10.255.19.77
set protocols bgp group int neighbor 10.255.19.79
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set routing-instances vpls instance-type vpls
set routing-instances vpls interface ge-0/2/2.0
set routing-instances vpls route-distinguisher 65000:1
set routing-instances vpls vrf-target target:65000:1
set routing-instances vpls protocols vpls site-range 3
set routing-instances vpls protocols vpls no-tunnel-services
set routing-instances vpls protocols vpls site asia site-identifier 1
set routing-instances vpls protocols vpls site asia interface ge-0/2/2.0
set routing-instances vpls protocols vpls vpls-id 100
```

PE2:

```
[edit]
set interfaces ge-0/0/7 gigether-options 802.3ad ae0
set interfaces ge-0/1/7 gigether-options 802.3ad ae0
set interfaces ge-0/2/3 gigether-options 802.3ad ae0
set interfaces ge-0/2/7 gigether-options 802.3ad ae0
set interfaces ge-2/0/0 gigether-options 802.3ad ae1
set interfaces ge-2/0/1 gigether-options 802.3ad ae1
set interfaces ge-2/0/2 gigether-options 802.3ad ae1
set interfaces ge-2/0/4 encapsulation ethernet-vpls
set interfaces ge-2/0/4 unit 0 family vpls
set interfaces ge-2/0/7 gigether-options 802.3ad ae0
set interfaces ge-2/0/9 unit 0 family inet address 1.1.1/30
set interfaces ge-2/0/9 unit 0 family mpls
set interfaces ge-2/1/7 gigether-options 802.3ad ae0
set interfaces ge-2/2/7 gigether-options 802.3ad ae0
set interfaces ge-2/3/7 gigether-options 802.3ad ae0
set interfaces ae0 unit 0 family inet address 11.11.11.2/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set interfaces ae1 unit 0 family inet address 10.1.1.1/30
set interfaces ae1 unit 0 family mpls
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path PE2-to-PE3 from 10.255.19.77
set protocols mpls label-switched-path PE2-to-PE3 to 10.255.19.79
set protocols mpls label-switched-path PE2-to-PE1 to 10.255.71.214
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.19.77
set protocols bgp group int family inet any
set protocols bgp group int family l2vpn signaling
set protocols bgp group int neighbor 10.255.71.214
set protocols bgp group int neighbor 10.255.19.79
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0
set protocols ospf area 0.0.0.0 interface ge-2/0/1.0
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols ospf area 0.0.0.0 interface ae1.0
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set routing-instances vpls instance-type vpls
set routing-instances vpls interface ge-2/0/4.0
set routing-instances vpls route-distinguisher 65000:1
set routing-instances vpls vrf-target target:65000:1
set routing-instances vpls protocols vpls site-range 3
set routing-instances vpls protocols vpls no-tunnel-services
set routing-instances vpls protocols vpls site 2 site-identifier 2
set routing-instances vpls protocols vpls site 2 interface ge-2/0/4.0
set routing-instances vpls protocols vpls vpls-id 100
```

**Step-by-Step Procedure** To configure PE1:

1. Configure PE1's interfaces.

```
[edit]
user@host# edit interfaces
[edit interfaces]
user@host# set ge-0/0/6 gigether-options 802.3ad ae0
user@host# set ge-0/1/6 gigether-options 802.3ad ae0
user@host# set ge-0/2/2 encapsulation ethernet-vpls
user@host# set ge-0/2/2 unit 0 family vpls
user@host# set ge-0/2/3 gigether-options 802.3ad ae0
user@host# set ge-0/2/6 gigether-options 802.3ad ae0
user@host# set ge-0/3/0 gigether-options 802.3ad ae0
user@host# set ge-0/3/1 gigether-options 802.3ad ae0
user@host# set ge-0/3/6 gigether-options 802.3ad ae0
user@host# set ge-1/0/6 gigether-options 802.3ad ae0
user@host# set ge-1/2/6 unit 0 family inet address 13.1.1.2/30
user@host# set ae0 unit 0 family inet address 11.11.11.1/30
user@host# set ae0 unit 0 family iso
user@host# set ae0 unit 0 family mpls
user@host# exit
```

2. On PE1, configure the packet header data to be used for per-flow load balancing.

```
[edit]
user@host# edit forwarding-options hash-key family multiservice
[edit forwarding-options hash-key family multiservice]
set source-mac
set destination-mac
set payload ip layer-3
user@host# exit
```

3. Configure the routing policy on PE1.

```
[edit]
user@host# edit policy-options
[edit policy-options]
user@host# set policy-statement exp-to-fwd term a from community grn-com
user@host# set policy-statement exp-to-fwd term a then install-nexthop lsp
PE1-to-PE2
user@host# set policy-statement exp-to-fwd term a then accept
user@host# set policy-options community grn-com members target:65000:1
user@host# exit
```

4. Configure PE1's routing protocols and MPLS.

```
[edit]
user@host# edit protocols
[edit protocols]
user@host# set rsvp interface all
user@host# set rsvp interface fxp0.0 disable
user@host# set mpls label-switched-path PE1-to-PE2 to 10.255.19.77
user@host# set mpls label-switched-path PE1-to-PE3 to 10.255.19.79
user@host# set mpls interface all
user@host# set mpls interface fxp0.0 disable
user@host# set bgp group int type internal
user@host# set bgp group int local-address 10.255.71.214
```

```

user@host# set bgp group int family inet any
user@host# set bgp group int family l2vpn signaling
user@host# set bgp group int neighbor 10.255.19.77
user@host# set bgp group int neighbor 10.255.19.79
user@host# set ospf traffic-engineering
user@host# set ospf area 0.0.0.0 interface all
user@host# set ospf area 0.0.0.0 interface fxp0.0 disable
user@host# exit

```

5. Configure VPLS on PE1.

```

[edit]
user@host# edit routing-instances vpls
[edit routing-instances vpls]
user@host# set instance-type vpls
user@host# set interface ge-0/2/2.0
user@host# set route-distinguisher 65000:1
user@host# set vrf-target target:65000:1
user@host# set protocols vpls site-range 3
user@host# set protocols vpls no-tunnel-services
user@host# set protocols vpls site asia site-identifier 1
user@host# set protocols vpls site asia interface ge-0/2/2.0
user@host# set protocols vpls vpls-id 100
user@host# exit

```

#### Step-by-Step Procedure

To configure PE2:

1. Configure PE2's interfaces.

```

[edit]
user@host# edit interfaces
[edit interfaces]
user@host# set ge-0/0/7 gigether-options 802.3ad ae0
user@host# set ge-0/1/7 gigether-options 802.3ad ae0
user@host# set ge-0/2/3 gigether-options 802.3ad ae0
user@host# set ge-0/2/7 gigether-options 802.3ad ae0
user@host# set ge-2/0/0 gigether-options 802.3ad ae1
user@host# set ge-2/0/1 gigether-options 802.3ad ae1
user@host# set ge-2/0/2 gigether-options 802.3ad ae1
user@host# set ge-2/0/4 encapsulation ethernet-vpls
user@host# set ge-2/0/4 unit 0 family vpls
user@host# set ge-2/0/7 gigether-options 802.3ad ae0
user@host# set ge-2/0/9 unit 0 family inet address 1.1.1.1/30
user@host# set ge-2/0/9 unit 0 family mpls
user@host# set ge-2/1/7 gigether-options 802.3ad ae0
user@host# set ge-2/2/7 gigether-options 802.3ad ae0
user@host# set ge-2/3/7 gigether-options 802.3ad ae0
user@host# set ae0 unit 0 family inet address 11.11.11.2/30
user@host# set ae0 unit 0 family iso
user@host# set ae0 unit 0 family mpls
user@host# set ae1 unit 0 family inet address 10.1.1.1/30
user@host# set ae1 unit 0 family mpls
user@host# exit

```

2. Configure PE1's routing protocols and MPLS.

```

[edit]

```



```

user@host# edit protocols
[edit protocols]
user@host# set rsvp interface all
user@host# set rsvp interface fxp0.0 disable
user@host# set mpls label-switched-path PE2-to-PE3 from 10.255.19.77
user@host# set mpls label-switched-path PE2-to-PE3 to 10.255.19.79
user@host# set mpls label-switched-path PE2-to-PE1 to 10.255.71.214
user@host# set mpls interface all
user@host# set mpls interface fxp0.0 disable
user@host# set bgp group int type internal
user@host# set bgp group int local-address 10.255.19.77
user@host# set bgp group int family inet any
user@host# set bgp group int family l2vpn signaling
user@host# set bgp group int neighbor 10.255.71.214
user@host# set bgp group int neighbor 10.255.19.79
user@host# set ospf traffic-engineering
user@host# set ospf area 0.0.0.0 interface lo0.0
user@host# set ospf area 0.0.0.0 interface ge-2/0/0.0
user@host# set ospf area 0.0.0.0 interface ge-2/0/1.0
user@host# set ospf area 0.0.0.0 interface ge-2/0/2.0
user@host# set ospf area 0.0.0.0 interface ae0.0
user@host# set ospf area 0.0.0.0 interface ae1.0
user@host# set ospf area 0.0.0.0 interface all
user@host# set ospf area 0.0.0.0 interface fxp0.0 disable
user@host# set ldp interface all
user@host# set ldp interface fxp0.0 disable
user@host# exit

```

### 3. Configure VPLS on PE1.

```

[edit]
user@host# edit routing-instances vpls
[edit routing-instances vpls]
user@host# set instance-type vpls
user@host# set interface ge-2/0/4.0
user@host# set route-distinguisher 65000:1
user@host# set vrf-target target:65000:1
user@host# set protocols vpls site-range 3
user@host# set protocols vpls no-tunnel-services
user@host# set protocols vpls site 2 site-identifier 2
user@host# set protocols vpls site 2 interface ge-2/0/4.0
user@host# set protocols vpls vpls-id 100
user@host# exit

```

## Results

From configuration mode, confirm your configuration by entering the **show forwarding-options**, **show interfaces**, **show policy-options**, **show protocols**, and **show routing-instances** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

PE1:

```

user@host# show forwarding-options
hash-key {

```

```
family multiservice {
  source-mac;
  destination-mac;
  payload {
    ip {
      layer-3;
    }
  }
}

user@host# show interfaces
ge-0/0/6 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-0/1/6 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-0/2/2 {
  encapsulation ethernet-vpls;
  unit 0 {
    family vpls;
  }
}
ge-0/2/3 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-0/2/6 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-0/3/0 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-0/3/1 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-0/3/6 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-1/0/6 {
  gigether-options {
    802.3ad ae0;
  }
}
```

```

    }
  }
  ge-1/2/6 {
    unit 0 {
      family inet {
        address 13.1.1.2/30;
      }
    }
  }
  ae0 {
    unit 0 {
      family inet {
        address 11.11.1.1/30;
      }
      family iso;
      family mpls;
    }
  }
}

user@host# show policy-options
policy-statement exp-to-fwd {
  term a {
    from community grn-com;
    then {
      install-nexthop lsp PE1-to-PE2;
      accept;
    }
  }
}
community grn-com members target:65000:1;

user@host# show protocols
mpls {
  label-switched-path PE1-to-PE2 {
    to 10.255.19.77;
  }
  label-switched-path PE1-to-PE3 {
    to 10.255.19.79;
  }
}
bgp {
  group int {
    type internal;
    local-address 10.255.71.214;
    family inet {
      any;
    }
    family l2vpn {
      signaling;
    }
    neighbor 10.255.19.77;
    neighbor 10.255.19.79;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {

```

```
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}

user@host# show routing-instances
vpls {
    instance-type vpls;
    interface ge-0/2/2.0;
    route-distinguisher 65000:1;
    vrf-target target:65000:1;
    protocols {
        vpls {
            site-range 3;
            no-tunnel-services;
            site asia {
                site-identifier 1;
                interface ge-0/2/2.0;
            }
            vpls-id 100;
        }
    }
}
```

PE2:

```
user@host# show interfaces
ge-0/0/7 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/1/7 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/2/3 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/2/7 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-2/0/0 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/0/1 {
    gigether-options {
        802.3ad ae1;
    }
}
```

```
    }  
  }  
  ge-2/0/2 {  
    gigether-options {  
      802.3ad ae1;  
    }  
  }  
  ge-2/0/4 {  
    encapsulation ethernet-vpls;  
    unit 0 {  
      family vpls;  
    }  
  }  
  ge-2/0/7 {  
    gigether-options {  
      802.3ad ae0;  
    }  
  }  
  ge-2/0/9 {  
    unit 0 {  
      family inet {  
        address 1.1.1.1/30;  
      }  
      family mpls;  
    }  
  }  
  ge-2/1/7 {  
    gigether-options {  
      802.3ad ae0;  
    }  
  }  
  ge-2/2/7 {  
    gigether-options {  
      802.3ad ae0;  
    }  
  }  
  ge-2/3/7 {  
    gigether-options {  
      802.3ad ae0;  
    }  
  }  
  ae0 {  
    unit 0 {  
      family inet {  
        address 11.11.11.2/30;  
      }  
      family iso;  
      family mpls;  
    }  
  }  
  ae1 {  
    unit 0 {  
      family inet {  
        address 10.1.1.1/30;  
      }  
      family mpls;  
    }  
  }
```

```
    }  
  }  
user@host# show protocols  
rsvp {  
  interface all;  
  interface fxp0.0 {  
    disable;  
  }  
}  
mpls {  
  label-switched-path PE2-to-PE3 {  
    from 10.255.19.77;  
    to 10.255.19.79;  
  }  
  label-switched-path PE2-to-PE1 {  
    to 10.255.71.214;  
  }  
  interface all;  
  interface fxp0.0 {  
    disable;  
  }  
}  
bgp {  
  group int {  
    type internal;  
    local-address 10.255.19.77;  
    family inet {  
      any;  
    }  
    family l2vpn {  
      signaling;  
    }  
    neighbor 10.255.71.214;  
    neighbor 10.255.19.79;  
  }  
}  
ospf {  
  traffic-engineering;  
  area 0.0.0.0 {  
    interface lo0.0;  
    interface ge-2/0/0.0;  
    interface ge-2/0/1.0;  
    interface ge-2/0/2.0;  
    interface ae0.0;  
    interface ae1.0;  
    interface all;  
    interface fxp0.0 {  
      disable;  
    }  
  }  
}  
ldp {  
  interface all;  
  interface fxp0.0 {  
    disable;  
  }  
}
```

```

    }
  }

user@host# show routing-instances
vpls {
  instance-type vpls;
  interface ge-2/0/4.0; ## 'ge-2/0/4.0' is not defined
  route-distinguisher 65000:1;
  vrf-target target:65000:1;
  protocols {
    vpls {
      site-range 3;
      no-tunnel-services;
      site 2 {
        site-identifier 2;
        interface ge-2/0/4.0;
      }
      vpls-id 100;
    }
  }
}

```

## Verification

You can monitor the operation of the routing instance by running the **show interfaces ae1.0 extensive** and **monitor interface traffic** commands.

For troubleshooting, you can configure tracing operations for all of the protocols.

### Related Documentation

- Configuring Point-to-Multipoint LSPs for an MBGP MVPN
- Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN
- show interfaces (Aggregated Ethernet)
- Multicast over Layer 3 VPNs





# Traffic Sampling Configuration

- [Traffic Sampling Configuration on page 211](#)
- [Minimum Traffic Sampling Configuration on page 212](#)
- [Configuring Traffic Sampling on page 213](#)
- [Disabling Traffic Sampling on page 215](#)
- [Configuring the Output File for Traffic Sampling on page 215](#)
- [Tracing Traffic-Sampling Operations on page 217](#)
- [Configuring Flow Aggregation \(cflowd\) on page 217](#)
- [Configuring Active Flow Monitoring Using Version 9 on page 220](#)
- [Traffic Sampling Examples on page 221](#)
- [Example: Sampling a Single SONET/SDH Interface on page 221](#)
- [Example: Sampling All Traffic from a Single IP Address on page 222](#)
- [Example: Sampling All FTP Traffic on page 223](#)

## Traffic Sampling Configuration

---

To configure traffic sampling, include the **sampling** statement at the **[edit forwarding-options]** hierarchy level:

```
[edit forwarding-options]
sampling {
  disable;
  family (inet | inet6 | mpls) {
    disable;
    output {
      aggregate-export-interval seconds;
      extension-service service-name;
      file {
        disable;
        filename filename;
        files number;
        size bytes;
        (stamp | no-stamp);
        (world-readable | no-world-readable);
      }
      flow-active-timeout seconds;
      flow-inactive-timeout seconds;
    }
  }
}
```

```
flow-server 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;
  version9 {
    template template-name;
  }
}
interface interface-name {
  engine-id number;
  engine-type number;
  source-address address;
}
}
input {
  max-packets-per-second number;
  maximum-packet-length bytes;
  rate number;
  run-length number;
}
traceoptions {
  file filename {
    files number;
    size bytes;
    (world-readable | no-world-readable);
  }
}
}
```

---

## Minimum Traffic Sampling 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 family-name]
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 {
    rate number;
  }
}

```

## Configuring Traffic Sampling

On routing platforms containing a Monitoring Services 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 {
  max-packets-per-second number;
  maximum-packet-length bytes;
  rate number;
  run-length number;
}

```

In Junos OS Release 8.3 and later, 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 maximum length of the sampled packet by using the **maximum-packet-length bytes** statement. For **bytes**, specify a value. For MX-Series devices with Modular Port Concentrators (MPCs), port-mirrored or sampled packets can be truncated (or clipped) to any length in the range of 1 to 255 bytes. Only 1 to 255 are valid values for packet truncation on these devices. For other devices, the range is from 0 to 9216. A maximum-packet-length value of zero represents that truncation is disabled, and the entire packet is mirrored or sampled.

Specify the sampling rate by setting the values for **rate** and **run-length** (see [Figure 17 on page 214](#)).

**Figure 17: Configure Sampling Rate**

### Rate and Run-length

Case #1 Rate =10, run-length =4



The forwarding plane provides support for random sampling that can be configured through the **rate** or **run-length** statement. The **rate** statement sets the ratio of the number of packets to be sampled on an average. For example, if you configure a rate of 10, on average every tenth packet (1 packet out of 10) 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.



**NOTE:** The **run-length** and **maximum-packet-length** configuration statements are not supported on MX80 routers.

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 217](#) and [“Configuring Active Flow Monitoring Using Version 9” on page 220](#).

The Junos OS 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.

**Related Documentation**

- Guidelines for Configuring Standard Firewall Filters
- Guidelines for Applying Standard Firewall Filters

## 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;
```

## Configuring the Output File for Traffic Sampling

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 family family-name 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 family (inet | inet6 | mpls)]** hierarchy level:

```
[edit forwarding-options sampling family (inet | inet6 | mpls) 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 family (inet | inet6 | mpls) output]
flow-inactive-timeout seconds;
```

To configure the interface that sends out monitored information, include the **interface** statement at the **[edit forwarding-options sampling output]** hierarchy level:

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output]
interface interface-name {
  engine-id number;
  engine-type number;
  source-address address;
}
```



**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 220](#).

## Traffic Sampling Output Format

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
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 family (inet | inet6 | mpls) 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 addr  port  port      len  num  frag  flags
# Feb  1 20:31:21
#              Dest      Src  Dest  Src Proto  TOS   Pkt  Intf   IP   TCP
#              addr      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>;
  no-remote-trace;
}
```

## 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 **flow-server** statement at the **[edit forwarding-options sampling output]** hierarchy level:

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output ]
flow-server 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 family (inet | inet6 | mpls) output 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 OS complies with Version 2.1b1 of cflowd. If you do not include the **caida-compliant** statement in the configuration, the Junos OS 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 family (inet | inet6 | mpls) output flow-server 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 the Output File for Traffic Sampling” on page 215](#).



**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

In Junos OS Release 8.3 and later, 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 a 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 Release 12.3 and the Flow Monitoring.

To enable the collection of traffic flows using the version 9 format, include the **version9** statement at the **[edit forwarding-options sampling family family-name output flow-server hostname]** hierarchy level:

```
[edit forwarding-options sampling family family-name output flow-server 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. In Junos OS Release 8.3 and later, you can configure sampling for 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. In Junos OS Release 9.5 and later, you can sample packets from both the **inet** and **mpls** protocol families at the same time. In Junos OS Release 10.4 and later, you can configure sampling for peer AS billing traffic for the **inet** and **ipv6** protocols only. For more information about how to configure traffic sampling, see [“Configuring Traffic Sampling” on page 213](#).

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 family family-name output flow-server 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 217](#).

- 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 the **[edit services monitoring]** hierarchy level. You also configure the traffic family **mpls** to sample MPLS packets.

```
[edit forwarding-options]
sampling {
  input {
    rate 1;
    run-length;
  }
  family inet {
    output {
      flow-server 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
      }
    }
  }
}
```

## Traffic Sampling Examples

The following sections provide examples of configuring traffic sampling:

- [Example: Sampling a Single SONET/SDH Interface on page 221](#)
- [Example: Sampling All Traffic from a Single IP Address on page 222](#)
- [Example: Sampling All FTP Traffic on page 223](#)

### Example: 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 {
    rate 100;
    run-length 2;
  }
  family inet {
    output {
      file {
        filename sonet-samples.txt;
        files 40;
        size 5m;
      }
    }
  }
}
```

---

## Example: 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 {
    rate 1;
  }
  family inet {
    output {
      file {
        filename samples-215-45-92-31.txt;
        files 100;
        size 100k;
      }
    }
  }
}

```

## Example: 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 {
    rate 10;
  }
  family inet {
    output {
      file {
        filename t3-ftp-traffic.txt;
        files 50;
        size 1m;
      }
    }
  }
}
```

## CHAPTER 12

# Traffic Forwarding and Monitoring Configuration

- [Configuring Traffic Forwarding and Monitoring on page 225](#)
- [Configuring Forwarding Table Filters on page 229](#)
- [Applying Filters to Forwarding Tables on page 231](#)
- [Configuring IPv6 Accounting on page 234](#)
- [Configuring Discard Accounting on page 234](#)
- [Configuring Flow Monitoring on page 235](#)
- [Forwarding Table Filters for Routing Instances on ACX Series Routers on page 236](#)
- [Configuring Load-Balance Groups on page 237](#)
- [Configuring Next-Hop Groups on page 237](#)
- [Configuring Per-Prefix Load Balancing on page 238](#)
- [Configuring Per-Flow Load Balancing Based on Hash Values on page 239](#)
- [Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239](#)
- [Configuring DNS and TFTP Packet Forwarding on page 241](#)
- [Preventing DHCP Spoofing on MX Series 3D Universal Edge Routers on page 245](#)
- [Configuring Port Mirroring on page 246](#)
- [Configuring Packet Capture on page 250](#)

## Configuring Traffic Forwarding and Monitoring

---

To configure forwarding options and traffic monitoring, 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-name {
    engine-id number;
    engine-type number;
    source-address address;
}
}
}
enhanced-hash-key {
    family [
        inet
            incoming-interface-index;
            no-destination-port;
            no-source-port;
            type-of-service;
        ]
        inet6 {
            incoming-interface-index;
            no-destination-port;
            no-source-port;
            traffic-class;
        }
        mpls {
            incoming-interface-index;
            label-1-exp;
            no-payload;
        }
        multiservice {
            incoming-interface-index;
            no-payload;
            outer-priority;
        }
    ]
    services-loadbalancing {
        family inet layer-3-services {
            incoming-interface-index;
            source-address;
        }
    }
}
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 {
      ether-pseudowire;
      ip {
        layer-3-only;
        port-data {
          source-msb;
          source-lsb;
          destination-msb;
          destination-lsb;
        }
      }
    }
  }
}
}
family multiservice {
  destination-mac;
  label-1;
  label-2;
  payload {
    ip {
      layer-3-only;
    }
  }
  source-mac;
}
}
helpers {
  bootp {
    client-response-ttl;
    description text-description;
    interface interface-group {
      client-response-ttl number;
      description text-description;
      maximum-hop-count number;
      minimum-wait-time seconds;
      no-listen;
      server address {
        logical-system logical-system-name <routing-instance [ <default>
          routing-instance-names ]>;
        routing-instance [ <default> routing-instance-names ];
      }
    }
  }
}

```

```

    }
    maximum-hop-count number;
    minimum-wait-time seconds;
    relay-agent-option;
    server [ addresses ];
}
domain {
    description text-description;
    server < [ routing-instance routing-instance-names ] >;
    interface interface-name {
        description text-description;
        no-listen;
        server < [ routing-instance routing-instance-names ] >;
    }
}
tftp {
    description text-description;
    server < [ routing-instance routing-instance-names ] >;
    interface interface-name {
        description text-description;
        no-listen;
        server < [ routing-instance routing-instance-names ] >;
    }
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size size>
        <world-readable | no-world readable>;
    flag flag;
    level severity-level;
    no-remote-trace;
}
}
load-balance {
    indexed-load-balance;
    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 {
                cflowd-collector;
            }
            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 bytes;
}
port-mirroring {
    family (ccc | inet | inet6 | vpls) {
        output {
            interface interface-name {
                next-hop address;
            }
            no-filter-check;
        }
        input {
            maximum-packet-length bytes;
            rate number;
            run-length number;
        }
    }
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size bytes>
        <world-readable | no-world-readable>;
    no-remote-trace;
}
}

```

## Configuring Forwarding Table Filters

The following sections describe the following topics:

- [Overview of Forwarding Table Filters on page 229](#)
- [Configuring a Forwarding Table Filter on page 230](#)

### Overview of Forwarding Table Filters

Forwarding table filters are defined the same as other firewall filters, but you apply them differently:

- Instead of applying forwarding table filters to interfaces, you apply them to forwarding tables, each of which is associated with a routing instance and a virtual private network (VPN).

- Instead of applying input and output 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 then performs a lookup for the forwarding table, thereby controlling which packets the router forwards on the interfaces.

When the router receives a packet, it determines the best route to the ultimate destination by looking in a forwarding table, which is associated with the VPN on which the packet is to be sent. The router then forwards the packet toward its destination through the appropriate interface.



**NOTE:** For transit packets exiting the router through the tunnel, forwarding table filtering is not supported on the interfaces you configure as the output interface for tunnel traffic.

## 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; it essentially controls which bearer packets the router accepts and forwards. To configure 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;
        }
      }
    }
  }
}
```

***family-name*** is the family address type: IPv4 (**inet**), IPv6 (**inet6**), Layer 2 traffic (**bridge**), or MPLS (**mpls**).

***term-name*** is a named structure in which match conditions and actions are defined.

***match-conditions*** are the criteria 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.

***action*** specifies what happens if a packet matches all criteria; 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.

**action-modifiers** 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.

To create a forwarding table, include the **instance-type** statement with the **forwarding** option at the **[edit routing-instances *instance-name*]** hierarchy level:

```
[edit]
routing-instances instance-name {
  instance-type forwarding;
}
```

To apply a forwarding table filter to a VPN routing and forwarding (VRF) table, include the **filter** and **input** statements at the **[edit routing-instance *instance-name* forwarding-options family *family-name*]** hierarchy level:

```
[edit routing-instances 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** and **input** statements 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** and **input** statements at the **[edit forwarding-options family inet]** hierarchy level:

```
[edit forwarding-options family inet]
filter {
  input filter-name;
}
```

#### Related Documentation

- Guidelines for Configuring Standard Firewall Filters
- Guidelines for Applying Standard Firewall Filters
- [Applying Filters to Forwarding Tables on page 231](#)

## Applying Filters to Forwarding Tables

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 apply a filter on the ingress

or egress packets of a forwarding table. You configure the filter at the **[edit firewall family *family-name*]** hierarchy level; for more information, see “[Configuring Forwarding Table Filters](#)” on page 229.

To apply a forwarding table filter on ingress packets of a forwarding table, include the **filter** and **input** statements 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** and **input** statements 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 OS Layer 2 Configuration Guide.

You can filter based upon destination-class information by applying a firewall filter on the egress packets 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 on egress packets of a forwarding table, include the **filter** and **output** statements 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 is applied on the ingress interface of the Flexible PIC Concentrator (FPC). If different packets to the same destination arrive on different FPCs, they might encounter different policers.



**NOTE:** You cannot simultaneously include the **interface-group** statement at the **[edit firewall family inet filter *filter-name* term *term-name* from]** hierarchy level and configure an egress forwarding table filter. 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.



**NOTE:** In Junos OS Release 8.4 and later, you can no longer configure this output statement for VPLS. You can continue to configure ingress forwarding table filters with the input statement at the [edit forwarding-options family vpls filter] hierarchy level.



**NOTE:** For T Series routers other than T4000, a packet forwarded by the forwarding table reaches the egress forwarding table filter irrespective of whether the packet is actually forwarded by the forwarding table or not. The packet reaches the egress filter even if the route points to reject or discard next hops.

On T4000 Type 5 FPC, the packet reaches the egress filter only if it is forwarded by the forwarding table.

To apply a forwarding table filter to a flood table, include the **flood** and **input** statements at the [edit forwarding-options family *family-name*] hierarchy level:

```
[edit forwarding-options family vpls]
flood {
  input filter-name;
}
```



**NOTE:** The **flood** statement is valid for the vpls protocol family only.

On MX Series 3D Universal Edge Routers, you can apply a forwarding table filter by using the **source-checking** statement at the [edit forwarding-options family inet6] hierarchy level:

```
[edit forwarding-options family inet6]
family inet6 {
  source-checking;
}
```

This discards IPv6 packets when the source address type is unspecified, loopback, multicast or link-local.

RFC 4291, *IP Version 6 Addressing Architecture*, refers to four address types that require special treatment when they are used as source addresses. The four address types are:

- Unspecified
- Loopback

- Multicast
- Link-Local Unicast

The loopback and multicast addresses must never be used as a source address in IPv6 packets. The unspecified and link-local addresses can be used as source addresses but routers must never forward packets that have these addresses as source addresses. Typically, packets that contain unspecified or link-local addresses as source addresses are delivered to the local host. If the destination is not the local host, then the packet must not be forwarded. Configuring this statement filters or discards IPv6 packets of these four address types.

---

## 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 remains enabled after a reboot of the router. To view IPv6 statistics, issue the **show interface statistics** operational mode command.

### Related Documentation

- Junos OS Operational Mode Commands

---

## 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 group group-name** 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 {
        engine-id number;
        engine-type number;
        source-address address;
    }
}
}

```

To configure the output flow aggregation, include the **cflowd** statement. For more information about flow aggregation, see [“Configuring Flow Aggregation \(cflowd\)” on page 217](#). 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 [Guidelines for Configuring Standard Firewall Filters](#) and [Guidelines for Applying Standard Firewall Filters](#).

You can use discard accounting for passive and active flow monitoring.

- Related Documentation**
- Flow Monitoring
  - Junos OS 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;
    }
    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 217](#). 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 [Guidelines for Configuring Standard Firewall Filters](#) and [Guidelines for Applying Standard Firewall Filters](#).

- Related Documentation**
- Flow Monitoring
  - Junos OS Class of Service Configuration Guide

## Forwarding Table Filters for Routing Instances on ACX Series Routers

Forwarding table filter is a mechanism by which all the packets forwarded by a certain forwarding table are subjected to filtering and if a packet matches the filter condition, the configured action is applied on the packet. You can use the forwarding table filter mechanism to apply a filter on all interfaces associated with a single routing instance with a simple configuration. You can apply a forwarding table filter to a routing instance of type forwarding and also to the default routing instance **inet.0**. To configure a forwarding table filter, include the **filter filter-name** statement at the **[edit firewall family inet]** hierarchy level.

The following limitations apply to forwarding table filters configured on routing instances:

- You cannot attach the same filter to more than one routing instance.
- You cannot attach the same filter at both the `[edit interfaces interface-name family inet filter input filter-name]` and `[edit routing-instances instance-name forwarding-options family inet filter input filter-name]` hierarchy level.
- You cannot attach a filter that is either interface-specific or a physical interface filter.
- You cannot attach a filter to the egress direction of routing instances.

**Related  
Documentation**

- [Configuring Forwarding Table Filters on page 229](#)

## Configuring Load-Balance Groups

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 Next-Hop Groups” on page 237](#).

## Configuring Next-Hop Groups

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.

**Related Documentation**

- [Configuring Port Mirroring on page 246](#)

---

## Configuring Per-Prefix Load Balancing

---

By default, the Junos OS 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 or switches share the same set of forwarding next hops for a given destination, they can elect the same forwarding next hop.

You can enable router-specific or switch-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-load-balance;
  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 (the default) through 65,535. If no hash seed is configured, the elected forwarding next hop is the same as in previous releases.

If you notice an issue with the load-balance distribution, try including the **indexed-load-balance** statement at the **[edit forwarding-options load-balance]** hierarchy level to see if this resolves the issue. The **indexed-load-balance** statement causes the

creation of a nexthop structure that is not a function of the hash only, but is also a function of the low-order bits of the IP address.



**CAUTION:** Including the **indexed-load-balance** statement causes an increase in memory usage on the device.

**indexed-load-balance;**

## Configuring Per-Flow Load Balancing Based on Hash Values

By default, the Junos OS uses a hashing method based only on the destination address to elect a forwarding next hop when multiple equal-cost paths are available. All Packet Forwarding Engine slots are assigned the same hash value by default.

You can enable router-specific or switch-specific load balancing by configuring the router or switch to assign a unique, load-balance hash value for each Packet Forwarding Engine slot.



**NOTE:** This feature is supported only on M120, M320, and MX Series routers.

To configure per-flow load balancing, include the **load-balance** statement at the **[edit forwarding-options]** hierarchy level:

```
[edit forwarding-options]
load-balance {
  indexed-load-balance;
  per-flow {
    hash-seed;
  }
}
```

To enable per-flow load balancing, you must include the **hash-seed** statement. The Junos OS automatically chooses a value for the hashing algorithm. You cannot configure a specific value for the **hash-seed** statement when you enable per-flow load balancing.

## Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

You can configure the router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) and 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, switch, or interface sees this broadcast message, it relays the message to a specified DHCP or BOOTP server.

You should configure the router, switch, or interface to be a DHCP and BOOTP relay agent if you have locally attached hosts and a distant DHCP or BOOTP server.

To configure the router or switch to act as a DHCP and 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 text-description;
  interface (interface-name | interface-group) {
    client-response-ttl number;
    description text-description;
    maximum-hop-count number;
    minimum-wait-time seconds;
    no-listen;
    server address {
      logical-system logical-system-name <routing-instance [ <default>
        routing-instance-names ]>;
      routing-instance [ <default> routing-instance-names ];
    }
  }
  maximum-hop-count number;
  minimum-wait-time seconds;
  relay-agent-option;
  server server-identifier {
    <logical-system logical-system-name>
    <routing-instance [ routing-instance-names ]>;
  }
}
```

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 or switch, 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 four 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 that specify the DHCP or BOOTP server for the router, switch, or interface, include the **server** statement. You can include multiple **server** statements.

To set an IP time-to-live (TTL) value for DHCP response packets sent to a DHCP client, include the **client-response-ttl** statement.

To use the DHCP relay agent option in relayed BOOTP/DHCP messages, include the **relay-agent-option** statement. This option is primarily useful for enabling DHCP forwarding

between different VRF routing instances. This option is documented in RFC 3046, *DHCP Relay Agent Information Option*.

You can also configure an individual logical interface to be a DHCP and BOOTP relay agent if you have locally attached hosts and a remote DHCP or BOOTP server connected to one of the router's or switch's interfaces. For more information, see the Junos OS System Basics Configuration Guide.

The following example demonstrates a BOOTP relay agent configuration.

```
user@host# show forwarding-options
helpers {
  bootp {
    description "dhcp relay agent global parameters";
    server 192.168.55.44;
    server 172.16.0.3 routing-instance c3;
    maximum-hop-count 10;
    minimum-wait-time 8;
    interface {
      fe-1/3/0 {
        description "use this info for this interface";
        server 10.10.10.10;
        server 192.168.14.14;
        maximum-hop-count 11;
        minimum-wait-time 3;
      }
      fe-1/3/1 {
        no-listen; ###ignore DHCPDISCOVER messages on this interface
      }
      all {
        description "globals apply to all other interfaces";
      }
    }
  }
}
```

## Configuring DNS and TFTP Packet Forwarding

You can configure the router or switch 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 or switch. The responding DNS or TFTP server recognizes the client address and sends a response directly to that address. By default, the router or switch 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 text-description;
    interface interface-name {
      description text-description;
      no-listen;
    }
  }
}
```

```
server [ addresses {  
    logical-system logical-system-name;  
    routing-instance instance-name;  
}  
}  
}  
tftp {  
    description text-description;  
    interface interface-name {  
        description text-description;  
        no-listen;  
        server address;  
        server logical-system name < [ routing-instance routing-instance-names ] >;  
        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 router or switch.

The following sections discuss the following:

- [Tracing BOOTP, DNS, and TFTP Forwarding Operations on page 242](#)
- [Example: Configuring DNS Packet Forwarding on page 244](#)

## 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 OS 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]
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag {
    address;
    all;
    config;
    domain;
    ifdb;
    io;
    main;
    port;
    rtsock;
    tftp;
    trace;
    ui;
    util;
  }
  level severity-level;
  no-remote-trace;
}
```

These statements are described in the following sections:

- [Configuring the Log Filename on page 243](#)
- [Configuring the Number and Size of Log Files on page 243](#)
- [Configuring Access to the Log File on page 244](#)
- [Configuring a Regular Expression for Lines to Be Logged on page 244](#)

### 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 filename` statement at the `[edit forwarding-options helpers traceoptions]` hierarchy level:

```
[edit forwarding-options helpers traceoptions]
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 traceoptions]` hierarchy level:

```
[edit forwarding-options helpers traceoptions]
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 **world-readable** option with the **file** statement at the **[edit forwarding-options helpers traceoptions]** hierarchy level:

```
[edit forwarding-options helpers traceoptions]
file world-readable;
```

To explicitly set the default behavior, include the **no-world-readable** option with the **file** statement at the **[edit forwarding-options helpers traceoptions]** hierarchy level:

```
[edit forwarding-options helpers traceoptions]
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** option with the **file** statement at the **[edit forwarding-options helpers traceoptions]** hierarchy level and specifying a regular expression (regex) to be matched:

```
[edit forwarding-options helpers traceoptions]
file filename match regular-expression;
```

## Example: Configuring DNS Packet Forwarding

Enable DNS packet request forwarding to all interfaces on a 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;
    }
  }
}
```

```

    }
  }
}

```

## Preventing DHCP Spoofing on MX Series 3D Universal Edge Routers

A problem that sometimes occurs with DHCP is *DHCP spoofing*, in which an untrusted client floods a network with DHCP messages. Often these attacks utilize source IP address spoofing to conceal the true source of the attack.

DHCP snooping helps prevent DHCP spoofing by copying DHCP messages to the control plane and using the information in the packets to create anti-spoofing filters. The anti-spoofing filters bind a client's MAC address to its DHCP-assigned IP address and use this information to filter spoofed DHCP messages. In a typical topology, a carrier edge router (in this function also referred to as the broadband services router [BSR]) connects the DHCP server and the MX Series router (or broadband services aggregator [BSA]) performing the snooping. The MX Series router connects to the client and the BSR.

DHCP snooping works as follows in the network topology mentioned above:

1. The client sends a DHCP discover message to obtain an IP address from the DHCP server.
2. The BSA intercepts the message and might add option 82 information specifying the slot, port, VPI/VCI, and so on.
3. The BSA then sends the DHCP discover message to the BSR, which converts it to a unicast packet and sends it to the DHCP server.
4. The DHCP server looks up the client's MAC address and option 82 information in its database. A valid client is assigned an IP address, which is returned to the client using a DHCP offer message. Both the BSR and BSA send this message upstream to the client.
5. The client examines the DHCP offer, and if it is acceptable, issues a DHCP request message that is sent to the DHCP server through the BSA and BSR.
6. The DHCP server confirms that the IP address is still available. If it is, the DHCP server updates its local tables and sends a DHCP ACK message to the client.
7. The BSR receives the DHCP ACK message and passes the message to the BSA.
8. The BSA creates an anti-spoofing filter by binding the IP address in the ACK message to the MAC address of the client. After this point, any DHCP messages from this IP address that are not bound to the client's MAC address are dropped.
9. The BSA sends the ACK message to the client so that the process of assigning a IP address can be completed.

You configure DHCP snooping by including within a DHCP group the appropriate interfaces of the BSA:

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name
  forwarding-options dhcp-relay group group-name]
```

**interface** *interface-name*;

In a VPLS environment, DHCP requests are forwarded over pseudowires. You can configure DHCP snooping over VPLS at the **[edit routing-instances *routing-instance-name*]** hierarchy level.

DHCP snooping works on a per learning bridge basis in bridge domains. Each learning domain must have an upstream interface configured. This interface acts as the flood port for DHCP requests coming from the client side. DHCP requests are be forwarded across learning domains in a bridge domain. You can configure DHCP snooping on bridge domains at the **[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*]** hierarchy level.

**Related Documentation**

- JUNOS® MX Series 3D Universal Edge Routers Solutions, Release 13.1

---

## Configuring Port Mirroring

---

*Port mirroring* is the ability of a router to send a copy of an IPv4 or IPv6 packet to an external host address or a packet analyzer for analysis. 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. A next-hop group can then be configured to forward copies of this duplicate packet to several interfaces. For more information about next-hop groups, see [“Configuring Next-Hop Groups” on page 237](#).

All M Series Multiservice Edge Routers, T Series Core Routers, and MX Series 3D Universal Edge Routers support port mirroring for IPv4 or IPv6. The M120, M320, and MX Series routers support port mirroring for IPv4 and IPv6 simultaneously.

Port mirroring for VPLS traffic is supported on M7i and M10i routers configured with an Enhanced CFEB (CFEB-E), on M120 routers, on M320 routers configured with an Enhanced III Flexible PIC Concentrators (FPCs), and MX Series routers.

In Junos OS Release 9.3 and later, port mirroring is supported for Layer 2 traffic on MX Series routers. For information about how to configure port mirroring for Layer 2 traffic, see the Junos OS Layer 2 Configuration Guide.

In Junos OS Release 9.6 and later, port mirroring is supported for Layer 2 VPN traffic on M120 routers and M320 routers configured with an Enhanced III FPC. You can also set the maximum length of the mirrored packet. When set, the mirrored packet is truncated to the specified length.

In the MPCs on M Series and MX Series routers, GRE and MPLS header information is not contained in the port-mirrored traffic corresponding to MPLS packets transmitted through IP-GRE tunnels.

## Configuration Guidelines

When configuring port mirroring, the following restrictions apply:

- Only transit data is supported.
- 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 can configure port mirroring for IPv4 and IPv6 simultaneously on the M120 and M320 routers and the MX Series routers.
- Ingress filtering of multicast packets is supported on all Dense Port Concentrators (DPCs) in MX Series routers. Egress filtering of multicast packets is supported for interfaces on Trio-based DPCs in MX Series routers only. Filtering of multicast packets based on destination address is not supported on M Series routers or T Series routers and is not supported for interfaces on I-chip ASIC-based DPCs in MX Series routers.
- By default, firewall filters cannot be applied to port-mirroring destination interfaces. To enable port-mirroring destination interfaces to support firewall filters, use the **no-filter-check** statement to disable filter checking on the interfaces. You can include the **no-filter-check** statement at the following hierarchy levels:
  - [edit forwarding-options port-mirroring family (inet | inet6 | ccc | vpls) output]
  - [edit forwarding-options port-mirroring instance *instance-name* family (inet | ccc | vpls) output]
- You must include a firewall filter with both the **accept** action and the **port-mirror** action modifier on the inbound interface. Port mirroring does not work if you specify the **discard** action.
- 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 **192.68.9.10** and the port-mirrored traffic is sent to **192.68.20.15** for analysis, the device associated with the latter address should not know a route to **192.68.9.10**. Also, it should not send the sampled packets back to the source address.
- On all routers except the MX Series router, 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. MX Series routers support more than one port-mirroring interface per router.
- You can configure multiple port-mirroring instances on the M120, M320, and MX Series routers.

- 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.
- 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.

## Configuring Port Mirroring

To configure port mirroring, include the **port-mirroring** statement at the **[edit forwarding-options]** hierarchy level:

```
[edit forwarding-options]
port-mirroring {
  family (ccc | inet | inet6 | vpls) {
    output {
      interface interface-name {
        next-hop address;
      }
      no-filter-check;
    }
    input {
      maximum-packet-length bytes;
      rate number;
      run-length number;
    }
  }
}
```

---

### Configuring the Port-Mirroring Address Family and Interface

To configure port mirroring, include the **port-mirroring** statement. To configure the address family type of traffic to sample, include the **family** statement. To configure the rate of sampling, length of sampling, and the maximum size for the mirrored packet, 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 determine whether 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 213.

## Configuring Multiple Port-Mirroring Instances

In Junos OS Release 9.5 and later, you can configure multiple port-mirroring instances on the M120, M320, and MX Series routers. On the M120 router, you can associate each instance with a specific Forwarding Engine Board (FEB). You cannot associate a port-mirroring instance with an FEB configured as a backup FEB. On the M320 router, you can associate each instance with a specific Flexible PIC Concentrator (FPC). Associating a port-mirroring instance with an FPC or an FEB enables you to mirror packets to different destinations. Multiple port-mirroring instances are also supported on MX Series routers. For information about configuring multiple port-mirroring instances on MX Series routers, see the Junos OS Layer 2 Configuration Guide.

To configure a port-mirroring instance, include the **instance** *port-mirroring-instance* statement at the **[edit forwarding-options port-mirroring]** hierarchy level:

```
[edit forwarding-options port-mirroring]
instance port-mirroring-instance-name {
  family (ccc | inet | inet6 | vpls) {
    output {
      interface interface-name {
        next-hop address;
      }
      no-filter-check;
    }
  }
  input {
    maximum-packet-length bytes;
    rate number;
    run-length number;
  }
}
```

### Configuring Port-Mirroring Instances

You can configure multiple port-mirroring instances. Specify a unique *port-mirroring-instance-name* for each instance you configure.

### Associating a Port-Mirroring Instance on M320 Routers

You can associate a port-mirroring instance with a specific FPC on an M320 router or with a specific FEB on an M120 router. You can associate only one port-mirroring instance with each FPC on an M320 router or with each FEB on an M120 router. On an M120 router, you cannot associate a port-mirroring instance with a FEB configured as a backup FEB.

To associate a port-mirroring instance with an FPC on an M320 router, include the **port-mirror-instance** *port-mirroring-instance-name* statement at the **[edit chassis fpc slot-number]** hierarchy level:

```
[edit chassis]
fpc slot-number {
  port-mirror-instance port-mirroring-instance-name;
}
```

For *slot-number*, specify the slot number of the FPC you want to associate with the port-mirroring instance. For *port-mirroring-instance-name*, specify the name of a port-mirroring instance you configured at the **[edit forwarding-options port-mirroring]** hierarchy level. For more information about configuring an FPC on an M320 router, see the Junos OS System Basics Configuration Guide.

### Associating a Port-Mirroring Instance on M120 Routers

To associate a port-mirroring instance with a FEB on an M120 router, include the **port-mirror-instance** *port-mirroring-instance-name* statement at the **[edit chassis feb slot-number]** hierarchy level:

```
[edit chassis]
feb slot-number {
```

```
port-mirror-instance port-mirroring-instance-name;  
}
```

For *slot-number*, specify the slot number of the FEB you want to associate with the port-mirroring instance. For *port-mirroring-instance-name*, specify the name of a port-mirroring instance you configured at the [edit forwarding-options port-mirroring] hierarchy level. For information about configuring FEB redundancy on an M120 router, see the Junos OS High Availability Configuration Guide. For information about configuring FPC-to-FEB connectivity on an M120 router, see the Junos OS System Basics Configuration Guide.

### Configuring MX Series 3D Universal Edge Routers and M120 Routers to Mirror Traffic Only Once

On MX Series and M120 routers only, you can configure port mirroring so that the router mirrors traffic only once. If you configure port mirroring on both ingress and egress interfaces, the same packet could be mirrored twice. To mirror packets only once and prevent the router from sending duplicate sampled packets to the same mirroring destination, include the **mirror-once** statement at the [edit forwarding-options port-mirroring] hierarchy level:

```
[edit forwarding-options port-mirroring]  
mirror-once;
```



**NOTE:** The **mirror-once** statement is supported only in the global port-mirroring instance.

## 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 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 bytes;  
}
```

To disable packet capture, include the **disable** statement. Packet capture is enabled by default.





**NOTE:** Packet capture can be configured only for the default routing instance.

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 the **var/tmp/** directory, the file is not re-created upon the next packet capture traffic on the interface. You must first disable and then enable PCAP functionality again to re-create 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:** A 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.

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.

## PART 3

# Administration

- [Routing Policy Reference on page 255](#)
- [Summary of Routing Policy Configuration Statements on page 263](#)
- [Summary of Traffic Sampling, Forwarding, and Monitoring Configuration Statements on page 297](#)



# Routing Policy Reference

- [Protocols That Can Be Imported to and Exported from the Routing Table on page 255](#)
- [Routing Tables Affected by Routing Policies on page 256](#)
- [Default Import and Export Policies for Protocols on page 257](#)
- [Routing Policy Match Conditions on page 258](#)
- [Protocol Support for Import and Export Policies on page 261](#)

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

**Table 22: 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
IS-IS	Yes	Yes
LDP	Yes	Yes
MPLS	Yes	No
OSPF	Yes	Yes
Protocol Independent Multicast (PIM) dense mode	Yes	Yes
PIM sparse mode	Yes	Yes
PIM sparse-dense mode	Yes	Yes

Table 22: Protocols That Can Be Imported to and Exported from the Routing Table (*continued*)

Protocol	Import	Export
Pseudoprotocol: <ul style="list-style-type: none"> <li>• Direct routes</li> <li>• Explicitly configured routes               <ul style="list-style-type: none"> <li>• Aggregate routes</li> <li>• Generated routes</li> <li>• Local routes</li> <li>• Static routes</li> </ul> </li> </ul>	Yes	No
Routing Information Protocol (RIP) and Routing Information Protocol next generation (RIPng)	Yes	Yes

## Routing Tables Affected by Routing Policies

Table 23 on page 256 lists the routing tables affected by default and user-defined routing policies and the types of routes that each routing table stores.

Table 23: Routing Tables Affected by Routing Policies

Routing Table	Type of Routes Stored
<b>inet.0</b>	Unicast IPv4 routes
<b>.inet.0 instance-name</b>	Unicast IPv4 routes for a particular routing instance
<b>inet.1</b>	Multicast IPv4 routes
<b>inet.2</b>	Unicast IPv4 routes for multicast reverse-path forwarding (RPF) lookup
<b>inet.3</b>	MPLS routes
<b>inet.flow.0</b>	Flow-specification routes
<b>instance-name.inetflow.0</b>	Flow-specification routes for Layer 3 VPNs
<b>mpls.0</b>	MPLS routes for label-switched path (LSP) next hops
<b>inet6.0</b>	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.

## Default Import and Export Policies for Protocols

Table 24: Default Import and Export Policies for Protocols

Importing or Exporting Protocol	Default Import Policy	Default Export Policy
BGP	Accept all received BGP IPv4 routes learned from configured neighbors and import into the <b>inet.0</b> routing table. Accept all received BGP IPv6 routes learned from configured neighbors and import into the <b>inet6.0</b> routing table.	Readvertise all learned BGP routes to all BGP speakers, while following protocol-specific rules that prohibit one IBGP speaker from readvertising routes learned from another IBGP speaker, unless it is functioning as a route reflector.
DVMRP	Accept all DVMRP routes and import into the <b>inet.1</b> routing table.	Accept and export active DVMRP routes.
IS-IS	Accept all IS-IS routes and import into the <b>inet.0</b> and <b>inet6.0</b> routing tables. (You cannot override or change this default policy.)	Reject everything. (The protocol uses flooding to announce local routes and any learned routes.)
LDP	Accept all LDP routes and import into the <b>inet.3</b> routing table.	Reject everything.
MPLS	Accept all MPLS routes and import into the <b>inet.3</b> routing table.	Accept and export active MPLS routes.
OSPF	Accept all OSPF routes and import into the <b>inet.0</b> 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 <b>inet.1</b> routing table.	Accept active PIM dense mode routes.
PIM sparse mode	Accept all PIM sparse mode routes and import into the <b>inet.1</b> routing table.	Accept and export active PIM sparse mode routes.
Pseudoprotocol: <ul style="list-style-type: none"> <li>• Direct routes</li> <li>• Explicitly configured routes: <ul style="list-style-type: none"> <li>• Aggregate routes</li> <li>• Generated routes</li> <li>• Static routes</li> </ul> </li> </ul>	Accept all direct and explicitly configured routes and import into the <b>inet.0</b> routing table.	<p>The pseudoprotocol cannot export any routes from the routing table because it is not a routing protocol.</p> <p>Routing protocols can export these or any routes from the routing table.</p>

Table 24: Default Import and Export Policies for Protocols (*continued*)

Importing or Exporting Protocol	Default Import Policy	Default Export Policy
RIP	Accept all RIP routes learned from configured neighbors and import into the <b>inet.0</b> 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 <b>inet6.0</b> 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 <a href="#">“Routing Policy Tests” on page 30</a> .	

## Routing Policy 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. 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 25 on page 258](#) 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 [“Configuring Match Conditions in Routing Policy Terms” on page 51](#).

Table 25: 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 25: 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 <a href="#">“How a Routing Policy Subroutine Is Evaluated”</a> on page 28.

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

- 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, all conditions must match the route for a match to occur.
- 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, all conditions must match the route for a match to occur.

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

Table 26 on page 260 summarizes key routing policy match conditions.

**Table 26: Summary of Key Routing Policy Match Conditions**

Match Condition	Description
<b>aggregate-contributor</b>	Matches routes that are contributing to a configured aggregate. This match condition can be used to suppress a contributor in an aggregate route.
<b>area <i>area-id</i></b>	Matches a route learned from the specified OSPF area during the exporting of OSPF routes into other protocols.
<b>as-path <i>name</i></b>	Matches the name of the path regular expression of an autonomous systems (AS). BGP routes whose AS path matches the regular expression are processed.
<b>color <i>preference</i></b>	Matches a color value. You can specify preference values that are finer-grained than those specified in the <b>preference</b> match conditions. The <b>color</b> value can be a number from 0 through 4,294,967,295 ( $2^{32} - 1$ ). A lower number indicates a more preferred route.
<b>community</b>	Matches the 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.)
<b>external [<i>type metric-type</i>]</b>	Matches external OSPF routes, including routes exported from one level to another. In this match condition, <b>type</b> is an optional keyword. The <b>metric-type</b> value can be either 1 or 2. When you do not specify <b>type</b> , this condition matches all external routes.
<b>interface <i>interface-name</i></b>	Matches the name or IP address of one or more router interfaces. Use this condition with protocols that are interface-specific. For example, do not use this condition with internal BGP (IBGP).  Depending on where the policy is applied, this match condition matches routes learned from or advertised through the specified interface.
<b>internal</b>	Matches a routing policy against the internal flag for simplified next-hop self policies.
<b>level <i>level</i></b>	Matches the IS-IS level. Routes that are from the specified level or are being advertised to the specified level are processed.
<b>local-preference <i>value</i></b>	Matches a BGP local preference attribute. The preference value can be from 0 through 4,294,967,295 ( $2^{32} - 1$ ).
<b>metric <i>metric</i></b> <b>metric2 <i>metric</i></b>	Matches a metric value. The <b>metric</b> value corresponds to the multiple exit discriminator (MED), and <b>metric2</b> corresponds to the IGP metric if the BGP next hop runs back through another route.
<b>neighbor <i>address</i></b>	Matches the address of one or more neighbors (peers).  For BGP export policies, the address can be for a directly connected or indirectly connected peer. For all other protocols, the address is for the neighbor from which the advertisement is received.
<b>next-hop <i>address</i></b>	Matches the 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.

Table 26: Summary of Key Routing Policy Match Conditions (*continued*)

Match Condition	Description
<b>origin value</b>	Matches the BGP origin attribute, which is the origin of the AS path information. The value can be one of the following: <ul style="list-style-type: none"> <li><b>egp</b>—Path information originated from another AS.</li> <li><b>igp</b>—Path information originated from within the local AS.</li> <li><b>incomplete</b>—Path information was learned by some other means.</li> </ul>
<b>preference preference</b> <b>preference2 preference</b>	Matches the preference value. You can specify a primary preference value ( <b>preference</b> ) and a secondary preference value ( <b>preference2</b> ). 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.
<b>protocol protocol</b>	Matches the 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: <b>aggregate</b> , <b>bgp</b> , <b>direct</b> , <b>dvmrp</b> , <b>isis</b> , <b>local</b> , <b>ospf</b> , <b>pim-dense</b> , <b>pim-sparse</b> , <b>rip</b> , <b>ripng</b> , or <b>static</b> .
<b>route-type value</b>	Matches the type of route. The value can be either <b>external</b> or <b>internal</b> .

## Protocol Support for Import and Export Policies

Table 27: Protocol Support for Import and Export Policies

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

Table 27: Protocol Support for Import and Export Policies (*continued*)

Protocol	Import Policy	Export Policy	Supported Levels
Pseudoprotocol—Explicitly configured routes, which include the following: <ul style="list-style-type: none"><li>Aggregate routes</li><li>Generated routes</li></ul>	Yes	No	Import: global
RIP and RIPng	Yes	Yes	Import: global, neighbor Export: group

## CHAPTER 14

# Summary of Routing Policy Configuration Statements

## aigp-originate

---

<b>Syntax</b>	<code>aigp-originate <i>distance</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> policy-options policy-statement <i>policy-name</i> term <i>term-name</i> then], [edit logical-systems <i>logical-system-name</i> policy-options policy-statement <i>policy-name</i> then], [edit policy-options policy-statement <i>policy-name</i> term <i>term-name</i> then], [edit policy-options policy-statement <i>policy-name</i> then]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	<p>Originate an accumulated interior gateway protocol (AIGP) BGP attribute for a given prefix by export policy, using the <b>aigp-originate</b> policy action.</p> <p>To originate an AIGP attribute, you need configure the policy action on only one node. The AIGP attribute is readadvertised if the neighbors are AIGP enabled with the <b>aigp</b> statement in the BGP configuration.</p>
<b>Default</b>	<p>If you omit the <b>aigp-originate</b> policy action, the node still readadvertises the AIGP BGP attribute if AIGP is enabled in the BGP configuration. However, the node does not originate its own AIGP attribute for local prefixes.</p> <p>As the route is readadvertised by downstream nodes, the cost of the AIGP attribute reflects the IGP distance to the prefix (zero + IGP distance or configured distance + IGP distance).</p>
<b>Options</b>	<p><b><i>distance</i></b>—(Optional) Associate an initial cost when advertising a local prefix with the AIGP BGP attribute.</p> <p><b>Range:</b> 0 through 4,294,967,295</p> <p><b>Default:</b> The initial cost associated with the AIGP attribute for a local prefix is zero. The <b><i>distance</i></b> option overrides the default zero value for the initial cost.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• Example: Configuring the Accumulated IGP Attribute for BGP</li><li>• aigp</li></ul>

## apply-path

---

<b>Syntax</b>	<code>apply-path path;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> <b>policy-options</b> <b>prefix-list</b> <i>name</i> ], [edit <b>policy-options</b> <b>prefix-list</b> <i>name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Expand a prefix list to include all prefixes pointed to by a defined path.
<b>Options</b>	<b>path</b> —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. You cannot add a path element, including wildcards, after a leaf statement. Path elements, including wildcards, can only be used after a container statement.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Prefix Lists on page 152</a></li></ul>

## as-path (Policy Options)

---

<b>Syntax</b>	<code>as-path name regular-expression;</code>
<b>Hierarchy Level</b>	[edit dynamic <a href="#">policy-options</a> ], [edit logical-systems <i>logical-system-name</i> <a href="#">policy-options</a> ], [edit <a href="#">policy-options</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Define an autonomous system (AS) path regular expression for use in a routing policy match condition.
<b>Options</b>	<b>name</b> —Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 65,536 characters long. To include spaces in the name, enclose it in quotation marks (" ").  <b>regular-expression</b> —One or more regular expressions used to match the AS path.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring AS Path Regular Expressions to Use as Routing Policy Match Conditions on page 133</a></li><li>• <a href="#">Configuring Routing Policies and Policy Objects in the Dynamic Database on page 77</a></li><li>• <a href="#">dynamic-db on page 272</a></li></ul>



## as-path-group

---

<b>Syntax</b>	<code>as-path-group <i>group-name</i> {     as-path <i>name</i> <i>regular-expression</i>; }</code>
<b>Hierarchy Level</b>	[edit dynamic <a href="#">policy-options</a> ], [edit logical-systems <i>logical-system-name</i> <a href="#">policy-options</a> ], [edit <a href="#">policy-options</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for dynamic database configuration introduced in Junos OS Release 9.5. Support for dynamic database configuration introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Define a group containing multiple AS path regular expressions for use in a routing policy match condition.
<b>Options</b>	<p><b><i>group-name</i></b>—Name that identifies the AS path group. One or more AS path regular expressions must be listed below the <b>as-path-group</b> hierarchy.</p> <p><b><i>name</i></b>—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><b><i>regular-expression</i></b>—One or more regular expressions used to match the AS path.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring AS Path Regular Expressions to Use as Routing Policy Match Conditions on page 133</a></li> <li>• <a href="#">Configuring Routing Policies and Policy Objects in the Dynamic Database on page 77</a></li> <li>• <a href="#">dynamic-db on page 272</a></li> </ul>

## community (Policy Options)

---

<b>Syntax</b>	<pre>community <i>name</i> {     invert-match;     members [ <i>community-ids</i> ]; }</pre>
<b>Hierarchy Level</b>	[edit dynamic <a href="#">policy-options</a> ], [edit logical-systems <i>logical-system-name</i> <a href="#">policy-options</a> ], [edit <a href="#">policy-options</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Define a community or extended community for use in a routing policy match condition.
<b>Options</b>	<p><b><i>name</i></b>—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 (" ").</p> <p><b><i>invert-match</i></b>—Invert the results of the community expression matching.</p> <p><b><i>members community-ids</i></b>—One or more community members. If you specify more than one member, you must enclose all members in brackets.</p> <p>The format for <b><i>community-ids</i></b> is:</p> <p style="padding-left: 40px;"><b><i>as-number:community-value</i></b></p> <p><b><i>as-number</i></b> is the AS number and can be a value in the range from 0 through 65,535. <b><i>community-value</i></b> is the community identifier and can be a number in the range from 0 through 65,535.</p> <p>You also can specify <b><i>community-ids</i></b> for communities as one of the following well-known community names, which are defined in RFC 1997, <i>BGP Communities Attribute</i>:</p> <ul style="list-style-type: none"><li>• <b><i>no-export</i></b>—Routes containing this community name are not advertised outside a BGP confederation boundary.</li><li>• <b><i>no-advertise</i></b>—Routes containing this community name are not advertised to other BGP peers.</li><li>• <b><i>no-export-subconfed</i></b>—Routes containing this community name are not advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation.</li></ul> <p>You can explicitly exclude BGP community information with a static route using the <b><i>none</i></b> option. Include <b><i>none</i></b> when configuring an individual route in the <b><i>route</i></b> portion of the <b><i>static</i></b> statement to override a <b><i>community</i></b> option specified in the <b><i>defaults</i></b> portion of the statement.</p>

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 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:** In Junos OS Release 9.1 and later, you can specify 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*, as well as the 2-byte AS numbers that are supported in earlier releases of the Junos OS. In plain-number format, you can configure a value in the range from 1 through 4,294,967,295. To configure a target or origin extended community that includes a 4-byte AS number in the plain-number format, append the letter “L” to the end of number. For example, a target community with the 4-byte AS number 334,324 and an assigned number of 132 is represented as **target:334324L:132**.

In Junos OS Release 9.2 and later, you can also use AS-dot notation when defining a 4-byte AS number for the target and origin extended communities. Specify two integers 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 represented in plain-number format as 65546 is represented in AS-dot notation as 1.10.

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

- Related Documentation**
- [Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions on page 31](#)
  - [Defining BGP Communities and Extended Communities for Use in Routing Policy Match Conditions on page 140](#)
  - [Configuring Routing Policies and Policy Objects in the Dynamic Database on page 77](#)
  - [dynamic-db on page 272](#)



---

## condition

---

<b>Syntax</b>	<code>condition <i>condition-name</i> {     if-route-exists <i>address</i> table <i>table-name</i>; }</code>
<b>Hierarchy Level</b>	[edit dynamic <a href="#">policy-options</a> ], [edit logical-systems <i>logical-system-name</i> <a href="#">policy-options</a> ], [edit <a href="#">policy-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Define a policy condition based on the existence of routes in specific tables for use in BGP export policies.
<b>Options</b>	<b>if-route-exists <i>address</i></b> —Specify the address of the route in question.  <b>table <i>table-name</i></b> —Specify a routing table.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Routing Policy Match Conditions Based on Routing Table Entries on page 174</a></li><li>• <a href="#">Configuring Routing Policies and Policy Objects in the Dynamic Database on page 77</a></li><li>• <a href="#">dynamic-db on page 272</a></li></ul>

## damping (Policy Options)

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> policy-options], [edit <i>policy-options</i>]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Define route flap damping properties to set on BGP routes.
<b>Options</b>	<p><b>disable</b>—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><b>half-life <i>minutes</i></b>—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.</p> <p><b>Range:</b> 1 through 45</p> <p><b>Default:</b> 15 minutes</p> <p> <b>NOTE:</b> For the half-life, configure a value that is less than the max-suppress. If you do not, the configuration is rejected.</p> <p><b>max-suppress <i>minutes</i></b>—Maximum hold-down time. <i>minutes</i> is the maximum time that a route can be suppressed no matter how unstable it has been.</p> <p><b>Range:</b> 1 through 720</p> <p><b>Default:</b> 60 minutes</p> <p> <b>NOTE:</b> For the max-suppress, configure a value that is greater than the half-life. If you do not, the configuration is rejected.</p> <p><b><i>name</i></b>—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><b>reuse <i>number</i></b>—Reuse threshold. <i>number</i> is the figure-of-merit value below which a suppressed route can be used again.</p> <p><b>Range:</b> 1 through 20,000</p> <p><b>Default:</b> 750 (unitless)</p>

**suppress *number***—Cutoff (suppression) threshold. *number* 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)

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

**Related Documentation**

- [Configuring BGP Flap Damping Parameters on page 181](#)

## dynamic-db

**Syntax** dynamic-db;

**Hierarchy Level** [edit logical-systems *logical-system-name* **policy-options as-path** *path-name*],  
[edit logical-systems *logical-system-name* **policy-options as-path-group** *group-name*],  
[edit logical-systems *logical-system-name* **policy-options community** *community-name*],  
[edit logical-systems *logical-system-name* **policy-options condition** *condition-name*],  
[edit logical-systems *logical-system-name* **policy-options policy-statement** *policy-statement-name*],  
[edit logical-systems *logical-system-name* **policy-options prefix-list** *prefix-list-name*],  
[edit **policy-options as-path** *path-name*],  
[edit **policy-options as-path-group** *group-name*],  
[edit **policy-options community** *community-name*],  
[edit **policy-options condition** *condition-name*],  
[edit **policy-options policy-statement** *policy-statement-name*],  
[edit **policy-options prefix-list** *prefix-list-name*]

**Release Information** Statement introduced in Junos OS Release 9.5.  
Statement introduced in Junos OS Release 9.5 for EX Series switches.

**Description** Define routing policies and policy objects that reference policies configured in the dynamic database at the **[edit dynamic]** hierarchy level.

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

**Related Documentation**


- [Configuring Routing Policies Based on Dynamic Database Configuration on page 77](#)

## export (Protocols BGP)

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Apply one or more policies to routes being exported from the routing table into BGP.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Understanding Route Advertisement</li> <li>Routing Policy Configuration Guide</li> <li><a href="#">import on page 282</a></li> </ul>

## export (Protocols IS-IS)

---

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	<p>Apply one or more policies to routes being exported from the routing table into IS-IS.</p> <p>All routing protocols store the routes that they learn in the routing table. The routing table uses this collected route information to determine the active routes to destinations. The routing table then installs the active routes into its forwarding table and exports them into the routing protocols. It is these exported routes that the protocols advertise.</p> <p>For each protocol, you control which routes the protocol stores in the routing table and which routes the routing table exports into the protocol from the routing table by defining a <i>routing policy</i> for that protocol.</p>
	<div><p><b>NOTE:</b> For IS-IS, you cannot apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol can easily lead to an inconsistent topology database.</p></div>
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• Example: Redistributing OSPF Routes into IS-IS</li><li>• Example: Configuring an IS-IS Default Route Policy on Logical Systems</li></ul>



## export (Protocols DVMRP)

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols dvmrp], [edit protocols dvmrp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply one or more policies to routes being exported from the routing table into DVMRP. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found, the routing table exports into DVMRP only the routes that it learned from DVMRP and direct routes.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">import on page 283</a></li> <li>• Example: Configuring DVMRP to Announce Unicast Routes</li> </ul>

## export (Protocols LDP)

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.
<b>Options</b>	<i>policy-names</i> —Name of one or more routing policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• Filtering Outbound LDP Label Bindings</li> </ul>

## export (Protocols MSDP)

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>],</p> <p>[edit protocols msdp],</p> <p>[edit protocols msdp group <i>group-name</i>],</p> <p>[edit protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit protocols msdp peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	Apply one or more policies to routes being exported from the routing table into MSDP.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Example: Configuring MSDP in a Routing Instance</li> <li><a href="#">import on page 284</a></li> </ul>

## export (Protocols OSPF)

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (ospf   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Apply one or more policies to routes being exported from the routing table into OSPF.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• Understanding OSPF Routing Policy</li> <li>• Import and Export Policies for Network Summaries Overview</li> <li>• <a href="#">import on page 285</a></li> <li>• Routing Policy Configuration Guide</li> </ul>

## export (Protocols PIM)

---


<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Apply one or more export policies to control outgoing PIM join and prune messages. PIM join and prune filters can be applied to PIM-SM and PIM-SSM messages. PIM join and prune filters cannot be applied to PIM-DM messages.
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>Filtering Outgoing PIM Join Messages</li></ul>

## export (Protocols PIM Bootstrap)

---

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim rp bootstrap family (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap family (inet   inet6)], [edit protocols pim rp bootstrap family (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap family (inet   inet6)]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Apply one or more export policies to control outgoing PIM bootstrap messages.
<b>Options</b>	<i>policy-names</i> —Name of one or more import policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>Configuring PIM Bootstrap Properties for IPv4</li><li>Configuring PIM Bootstrap Properties for IPv4 or IPv6</li><li><a href="#">import (Protocols PIM Bootstrap) on page 287</a></li></ul>

## export (Protocols RIP)

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i>],</p> <p>[edit protocols rip group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Apply a policy to routes being exported to the neighbors.</p> <p>By default, RIP does not export routes it has learned to its neighbors. To enable RIP to export routes, apply one or more export policies.</p> <p>If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the <b>metric-in</b> and <b>metric-out</b> statements.</p>
	<div>  <p><b>NOTE:</b> The export policy on RIP does not support manipulating routing information of the next hop.</p> </div>
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Example: Configuring RIP</li> <li><a href="#">import on page 288</a></li> </ul>

## export (Protocols RIPng)

---

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ripng group <i>group-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> ], [edit protocols ripng group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for routing instances introduced in Junos OS Release 9.0.
<b>Description</b>	<p>Apply a policy or list of policies to routes being exported to the neighbors.</p> <p>By default, RIPng does not export routes it has learned to its neighbors. To have RIPng export routes, apply one or more export policies. To apply export policies and to filter routes being exported from the local routing device to its neighbors, include the <b>export</b> statement and list the name of the policy to be evaluated.</p> <p>You can define one or more export policies. If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the <b>metric-in</b> and <b>metric-out</b> statements.</p>
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• Example: Configuring RIPng</li><li>• <a href="#">import on page 289</a></li></ul>

## export (Routing Options)

---

<b>Syntax</b>	<code>export [ <i>policy-name</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options forwarding-table], [edit logical-systems <i>logical-system-name</i> routing-options forwarding-table], [edit routing-instances <i>routing-instance-name</i> routing-options forwarding-table], [edit routing-options forwarding-table]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	Apply one or more policies to routes being exported from the routing table into the forwarding table.
<b>Options</b>	<i>policy-name</i> —Name of one or more policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• Example: Load Balancing BGP Traffic</li><li>• Routing Policy Configuration Guide</li></ul>

## import (Protocols BGP)

---

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>bgp group <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>bgp group <i>group-name</i> neighbor <i>address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp group <i>group-name</i>],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i></code> <code>neighbor <i>address</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Apply one or more routing policies to routes being imported into the Junos OS routing table from BGP.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• Example: Configuring BGP Interactions with IGP</li><li>• Understanding Route Advertisement</li><li>• <a href="#">Importing and Exporting Routes on page 17</a></li><li>• Routing Policy Configuration Guide</li><li>• <a href="#">export on page 273</a></li></ul>



## import (Protocols DVMRP)

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols dvmrp], [edit protocols dvmrp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply one or more policies to routes being imported into the routing table from DVMRP. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found, DVMRP shares with the routing table only those routes that were learned from DVMRP routers.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">export on page 275</a></li> <li>• Example: Configuring DVMRP to Announce Unicast Routes</li> </ul>

## import (Protocols LDP)

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.
<b>Options</b>	<i>policy-names</i> —Name of one or more routing policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• Filtering Inbound LDP Label Bindings</li> </ul>

## import (Protocols MSDP)

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<pre> [edit logical-systems <i>logical-system-name</i> protocols msdp], [edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols msdp peer <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>], [edit protocols msdp], [edit protocols msdp group <i>group-name</i>], [edit protocols msdp group <i>group-name</i> peer <i>address</i>], [edit protocols msdp peer <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols msdp], [edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols msdp group <i>group-name</i> peer <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols msdp peer <i>address</i>] </pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	Apply one or more policies to routes being imported into the routing table from MSDP.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Example: Configuring MSDP in a Routing Instance</li> <li><a href="#">export on page 276</a></li> </ul>

## import (Protocols OSPF)

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (ospf   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Filter OSPF routes from being added to the routing table.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Understanding OSPF Routing Policy</li> <li>Import and Export Policies for Network Summaries Overview</li> <li><a href="#">export on page 277</a></li> <li>Routing Policy Configuration Guide</li> </ul>

## import (Protocols PIM)

---

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

## import (Protocols PIM Bootstrap)

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bootstrap (inet   inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap (inet   inet6)],</p> <p>[edit protocols pim rp bootstrap (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp bootstrap (inet   inet6)]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Apply one or more import policies to control incoming PIM bootstrap messages.
<b>Options</b>	<i>policy-names</i> —Name of one or more import policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Configuring PIM Bootstrap Properties for IPv4</li> <li>Configuring PIM Bootstrap Properties for IPv4 or IPv6</li> <li><a href="#">export (Protocols PIM Bootstrap) on page 278</a></li> </ul>

## import (Protocols RIP)

---

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols rip],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> neighbor</code> <code>  <i>neighbor-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>  rip],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>  rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit protocols rip],</code> <code>[edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols rip],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> neighbor</code> <code>  <i>neighbor-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Apply one or more policies to routes being imported by the local routing device from neighbors.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• Example: Applying Policies to RIP Routes Imported from Neighbors</li><li>• Routing Policy Configuration Guide</li><li>• <a href="#">export on page 279</a></li></ul>

## import (Protocols RIPng)

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols ripng],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols ripng],</p> <p>[edit protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>
<b>Description</b>	Apply one or more policies to routes being imported into the local routing device from its neighbors.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• Example: Applying Policies to RIPng Routes Imported from Neighbors</li> <li>• <a href="#">export on page 280</a></li> </ul>

## import (Routing Options)

---

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution rib], [edit logical-systems <i>logical-system-name</i> routing-options resolution rib], [edit routing-instances <i>routing-instance-name</i> routing-options resolution rib], [edit routing-options resolution rib]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Specify one or more import policies to use for route resolution.
<b>Options</b>	<i>policy-names</i> —Name of one or more import policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Route Resolution on PE Routers</a></li></ul>

## policy-options

---

<b>Syntax</b>	<code>policy-options { ... }</code>
<b>Hierarchy Level</b>	[edit], [edit dynamic], [edit dynamic-profiles <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Support at the [edit dynamic-profiles] hierarchy level introduced in Junos OS Release 11.4.
<b>Description</b>	Configure routing policy.  The statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Defining Routing Policies on page 50</a></li></ul>



## policy-statement

<b>Syntax</b>	<pre> policy-statement <i>policy-name</i> {   term <i>term-name</i> {     from {       family <i>family-name</i>;       match-conditions;       policy <i>subroutine-policy-name</i>;       prefix-list <i>prefix-list-name</i>;       prefix-list-filter <i>prefix-list-name</i> match-type &lt;actions&gt;;       route-filter <i>destination-prefix</i> match-type &lt;actions&gt;;       source-address-filter <i>source-prefix</i> match-type &lt;actions&gt;;     }     to {       match-conditions;       policy <i>subroutine-policy-name</i>;     }     then <i>actions</i>;   } } </pre>
<b>Hierarchy Level</b>	<p>[edit dynamic <a href="#">policy-options</a>],  [edit logical-systems <i>logical-system-name</i> <a href="#">policy-options</a>],  [edit <a href="#">policy-options</a>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.  Statement introduced in Junos OS Release 9.0 for EX Series switches.  Support for configuration in the dynamic database introduced in Junos OS Release 9.5.  Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.  <b>inet-mdt</b> option introduced in Junos OS Release 10.0R2.  Statement introduced in Junos OS Release 11.3 for the QFX Series.  <b>route-target</b> option introduced in Junos OS Release 12.2.</p>
<b>Description</b>	<p>Define a routing policy, including subroutine policies.</p> <p>To list the routing policies under the <b>[edit policy-options]</b> hierarchy level by <b>policy-statement <i>policy-name</i></b> in alphabetical order, enter the <b>show policy-options</b> configuration command.</p>
<b>Options</b>	<p><b>actions</b>—(Optional) One or more actions to take if the conditions match. The actions are described in “<a href="#">Configuring Flow Control Actions</a>” on page 60.</p> <p><b>family <i>family-name</i></b>—(Optional) Specify an address family protocol. Specify <b>inet</b> for IPv4. Specify <b>inet6</b> for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify <b>iso</b>. For IPv4 multicast VPN traffic, specify <b>inet-mvpn</b>. For IPv6 multicast VPN traffic, specify <b>inet6-mvpn</b>. For multicast-distribution-tree (MDT) IPv4 traffic, specify <b>inet-mdt</b>. For BGP route target VPN traffic, specify <b>route-target</b>.</p>



**NOTE:** When family is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the family statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

**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 “[Configuring Match Conditions in Routing Policy Terms](#)” on page 51.

**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 (“ ”). Policy names cannot take the form `__.*-internal__`, as this form is reserved. For information about how to configure subroutines, see “[Configuring Subroutines in Routing Policy Match Conditions](#)” on page 170.

**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.

**prefix-list-filter prefix-list-name**—Name of a prefix list to evaluate using qualifiers; **match-type** is the type of match (see “[Configuring Prefix List Filters](#)” on page 155), and **actions** is the action to take if the prefixes match.

**route-filter destination-prefix match-type <actions>**—(Optional) List of routes on which to perform an immediate match; **destination-prefix** is the IPv4 or IPv6 route prefix to match, **match-type** is the type of match (see “[Configuring Route Lists](#)” on page 156), and **actions** is the action to take if the **destination-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 “[Configuring Route Lists](#)” on page 156), 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 [“Configuring Flow Control Actions” on page 60](#) and [“Configuring Actions That Manipulate Route Characteristics” on page 60](#).

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

**Related Documentation**

- [Defining Routing Policies on page 50](#)
- [Configuring Routing Policies and Policy Objects in the Dynamic Database on page 77](#)
- [dynamic-db on page 272](#)

## prefix-list

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

**Hierarchy Level** [edit dynamic [policy-options](#)],  
[edit logical-systems *logical-system-name* [policy-options](#)],  
[edit [policy-options](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Support for configuration in the dynamic database introduced in Junos OS Release 9.5.  
Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.  
Support for the **vpls** protocol family introduced in Junos OS Release 10.2.

**Description** Define a list of IPv4 or IPv6 address prefixes for use in a routing policy statement or firewall filter statement.

You can configure up to 85,325 prefixes in each prefix list. To configure more than 85,325 prefixes, configure multiple prefix lists and apply them to multiple firewall filter terms.

**Options** *name*—Name that identifies the list of IPv4 or IPv6 address prefixes.

*ip-addresses*—List of IPv4 or IPv6 address prefixes, one IP address per line in the configuration.

The remaining statement is explained separately.

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

**Related Documentation**

- [Configuring Prefix Lists for Use in Routing Policy Match Conditions on page 151](#)
- Firewall Filter Match Conditions Based on Address Fields
- Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List

## prefix-list-filter

---

<b>Syntax</b>	<code>prefix-list-filter <i>prefix-list-name</i> <i>match-type</i> &lt;<i>actions</i>&gt;;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> <a href="#">policy-options</a> ], [edit <a href="#">policy-options</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Evaluate a list of prefixes within a prefix list using specified qualifiers.
<b>Options</b>	<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>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Prefix Lists for Use in Routing Policy Match Conditions on page 151</a></li></ul>

## rtf-prefix-list

<b>Syntax</b>	<code>rtf-prefix-list name route-targets</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> <a href="#">policy-options</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> <a href="#">policy-options</a> <a href="#">policy-statement</a> <i>policy-name</i> term <i>term-name</i>],</p> <p>[edit <a href="#">policy-options</a>],</p> <p>[edit <a href="#">policy-options</a> <a href="#">policy-statement</a> <i>policy-name</i> term <i>term-name</i>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	<p>Define a list of route target prefixes for use in a routing policy statement. These prefixes are only useful for filtering routes in the <code>bpg.target.0</code> table.</p> <p>The route target filtering prefix is in the format: <i>AS number:route target extended community/length</i>. The first number represents the autonomous system (AS) of the device that sent the advertisement. The second group of numbers represent the route target extended community. The format of the extended community is the same as the extended community type <b>target</b>. For more information about extended communities, see <a href="#">“Defining BGP Communities and Extended Communities for Use in Routing Policy Match Conditions” on page 140</a>.</p> <p>In this route target prefix example 200:200:101/96, 200 is the AS number, 200:101 is the BGP extended community used for the route target, and 96 is the prefix length.</p> <p>For more information about the route target community, see RFC 4360, <i>BGP Extended Communities Attribute</i>.</p> <p>For more information about the route target filtering prefix format, see RFC 4684, <i>Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)</i>.</p>
<b>Options</b>	<p><b>name</b>—Name that identifies the list of route target filtering prefixes. 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 ( “ ” ).</p> <p><b>route-targets</b>—List of route target filtering prefixes, one route target filter per line in the configuration. When you use the <b>rtf-prefix-list</b> statement as a match condition, you do not have the option of configuring the list of route target filtering prefixes. You must first define and configure the route target filtering prefixes with the <a href="#">policy-options</a> statement.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• Example: Configuring an Export Policy for BGP Route Target Filtering</li> <li>• Configuring BGP Route Target Filtering for VPNs</li> <li>• Understanding Proxy BGP Route Target Filtering</li> </ul>

- [Defining BGP Communities and Extended Communities for Use in Routing Policy Match Conditions on page 140](#) in the Routing Policy Configuration Guide
- family route-target

## CHAPTER 15

# Summary of Traffic Sampling, Forwarding, and Monitoring Configuration Statements

## accounting

```
Syntax  accounting group-name {
        output {
            aggregate-export-interval seconds;
            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-name {
                engine-id number;
                engine-type number;
                source-address address;
            }
        }
    }
```

**Hierarchy Level** [edit [forwarding-options](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** Specify discard accounting instance name and options.

The statements are explained separately.

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

**Related Documentation**

- [Configuring Discard Accounting on page 234](#)



## aggregation

<b>Syntax</b>	<pre>aggregation {   autonomous-system;   destination-prefix;   protocol-port;   source-destination-prefix {     caida-compliant;   }   source-prefix; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options accounting output hostname</a> ], [edit <a href="#">forwarding-options sampling family</a> (inet   inet6   mpls) <a href="#">output flow-server hostname</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<p><b>autonomous-system</b>—Aggregate by autonomous system (AS) number.</p> <p><b>caida-compliant</b>—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 OS records source and destination mask length values in compliance with the <i>cflowd Configuration Guide</i>, dated August 30, 1999.</p> <p><b>destination-prefix</b>—Aggregate by destination prefix.</p> <p><b>protocol-port</b>—Aggregate by protocol and port number.</p> <p><b>source-destination-prefix</b>—Aggregate by source and destination prefix.</p> <p><b>source-prefix</b>—Aggregate by source prefix.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Flow Aggregation (cflowd) on page 217</a></li> </ul>

## autonomous-system-type

---

<b>Syntax</b>	<code>autonomous-system-type (origin   peer);</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options accounting</a> <code>output cflowd hostname</code> ], [edit <a href="#">forwarding-options sampling family</a> (inet   inet6   mpls) <code>output flow-server hostname</code> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the type of AS numbers that cflowd exports.
<b>Default</b>	<code>origin</code>
<b>Options</b>	<b>origin</b> —Export origin AS numbers of the packet source address in the Source Autonomous System cflowd field.  <b>peer</b> —Export peer AS numbers through which the packet passed in the Source Autonomous System cflowd field.
<b>Required Privilege Level</b>	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flow Aggregation (cflowd) on page 217</a></li></ul>

## bootp

<b>Syntax</b>	<pre> bootp {   client-response-ttl <i>number</i>;   description <i>text-description</i>;   apply-secondary-as-giaddr <i>text-description</i>; (QFX platforms only)   interface (<i>interface-name</i>   <i>interface-group</i>) {     client-response-ttl <i>number</i>;     description <i>text-description</i>;     maximum-hop-count <i>number</i>;     minimum-wait-time <i>seconds</i>;     no-listen;     server address {       logical-system <i>logical-system-name</i> &lt;routing-instance [ &lt;default&gt;         <i>routing-instance-names</i> ]&gt;;       routing-instance [ &lt;default&gt; <i>routing-instance-names</i> ];     }     smart-relay-agent (QFX platforms only)   }   maximum-hop-count <i>number</i>;   minimum-wait-time <i>seconds</i>;   relay-agent-option;   server address {     &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance       [ <i>routing-instance-names</i> ]&gt;;   } } </pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpers</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p>
<b>Description</b>	<p>Configures a router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent.</p> <p>DHCP relaying is disabled.</p>
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239</a></li> <li>• <a href="#">Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)</a></li> </ul>

## cflowd (Discard Accounting)

---

Syntax	<pre>cflowd hostname {   aggregation {     autonomous-system;     destination-prefix;     protocol-port;     source-destination-prefix {       caida-compliant;     }     source-prefix;   }   autonomous-system-type (origin   peer);   port port-number;   version format; }</pre>
Hierarchy Level	[edit <a href="#">forwarding-options accounting group-name output</a> ]
Release Information	Statement introduced before Junos OS 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 cfdcollect.</p> <p>You can configure up to one version 5 and one version 8 flow format at the [edit <a href="#">forwarding-options accounting group-name output</a>] hierarchy level.</p>
Options	<p><b>hostname</b>—The IP address or identifier of the host system (the workstation running the cflowd utility).</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flow Aggregation (cflowd) on page 217</a></li></ul>

## cflowd (Flow Monitoring)

<b>Syntax</b>	<code>cflowd hostname {     port port-number; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options monitoring group-name</a> family inet output]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect.</p> <p>You can configure up to eight version 5 flow formats at the [edit <a href="#">forwarding-options monitoring group-name output</a>] hierarchy level. Version 8 flow formats are not supported for flow-monitoring applications.</p>
<b>Options</b>	<p><b>hostname</b>—The IP address or identifier of the host system (the workstation running the cflowd utility).</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Flow Monitoring on page 235</a></li> </ul>

## client-response-ttl

<b>Syntax</b>	<code>client-response-ttl number;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpers bootp</a> ], [edit <a href="#">forwarding-options helpers bootp interface</a> (interface-name   interface-group)]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p>
<b>Description</b>	Set the IP time-to-live (TTL) value in DHCP response packets sent to a DHCP client.
<b>Options</b>	<p><b>number</b>—Decrement amount.</p> <p><b>Default:</b> None</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239</a></li> </ul>

## description (Forwarding Options)

---

<b>Syntax</b>	<code>description <i>text-description</i>;</code>
<b>Hierarchy Level</b>	<code>[edit forwarding-options helpers bootp]</code> , <code>[edit forwarding-options helpers bootpinterface (<i>interface-name</i>   <i>interface-group</i>)]</code> , <code>[edit forwarding-options helpers domain]</code> , <code>[edit forwarding-options helpers domain interface <i>interface-name</i>]</code> , <code>[edit forwarding-options helpers tftp]</code> , <code>[edit forwarding-options helpers tftpinterface <i>interface-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Describe a BOOTP, DHCP, Domain Name System (DNS), or Trivial File Transfer Protocol (TFTP) service, or an interface that is configured for the service.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring DNS and TFTP Packet Forwarding on page 241</a></li><li>• <a href="#">Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239</a></li></ul>

## dhcp-relay (DHCP Spoofing Prevention)

<b>Syntax</b>	<pre>dhcp-relay {   group group-name {     interface interface-name;   } }</pre>
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> forwarding-options], [edit routing-instances <i>routing-instance-name</i> forwarding-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4 (MX Series routers only).
<b>Description</b>	<p>Configure Dynamic Host Configuration Protocol (DHCP) snooping on the router. When acting as a snooping agent, the MX Series router typically is located between the client and the DHCP relay agent. It creates filters by “snooping” DHCP messages and binding DHCP-issued IP addresses to the MAC address of the client. These filters help prevent DHCP spoofing.</p> <p>Configure DHCP snooping by including the appropriate interfaces in the DHCP relay configuration.</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Preventing DHCP Spoofing on MX Series 3D Universal Edge Routers on page 245</a></li> </ul>

## disable (Forwarding Options)

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	[edit forwarding-options packet-capture], [edit forwarding-options sampling], [edit forwarding-options sampling family <i>family-name</i> output file]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Supported added at the [edit forwarding-options packet-capture] hierarchy level on J Series Services Routers in Junos OS Release 7.5.
<b>Description</b>	Disable traffic sampling or (on J Series Services Routers) packet capture.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Packet Capture on page 250</a></li> <li>• <a href="#">Disabling Traffic Sampling on page 215</a></li> </ul>

## domain

---

Syntax	<pre>domain {   <b>description</b> <i>text-description</i>;   <b>interface</b> <i>interface-name</i> {     broadcast;     <b>description</b> <i>text-description</i>;     <b>no-listen</b>;     <b>server</b> <i>address</i> &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance       <i>routing-instance-name</i>&gt;;   }   <b>server</b> <i>address</i> &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance     <i>routing-instance-name</i>&gt;; }</pre>
Hierarchy Level	[edit <a href="#">forwarding-options helpers</a> ]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Enable DNS request packet forwarding.  The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring DNS and TFTP Packet Forwarding on page 241</a></li></ul>

## export-format

---

Syntax	<pre>export-format cflowd-version-5;</pre>
Hierarchy Level	[edit <a href="#">forwarding-options monitoring</a> <i>group-name</i> <a href="#">family</a> inet <a href="#">output</a> ]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Flow monitoring export format.
Options	<b>cflowd-version-5</b> —cflowd version 5.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flow Monitoring on page 235</a></li></ul>



## enhanced-hash-key

<b>Syntax</b>	<pre> enhanced-hash-key {   family {     inet       incoming-interface-index;       no-destination-port;       no-source-port;       type-of-service;     }     inet6 {       incoming-interface-index;       no-destination-port;       no-source-port;       traffic-class;     }     mpls {       incoming-interface-index;       label-1-exp;       no-payload;     }     multiservice {       incoming-interface-index;       no-payload;       outer-priority;     }     services-loadbalancing {       family inet {         layer-3-services {           destination-address           incoming-interface-index;           source-address;         }       }       family inet6 {         layer-3-services {           destination-address           incoming-interface-index;           source-address;         }       }     }   } } </pre>
<b>Hierarchy Level</b>	<p>[edit forwarding-options],  [edit logical-systems logical-system-name routing-instances instance-name forwarding-options],  [edit routing-instances instance-name forwarding-options]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.1.  <b>services-loadbalancing</b> option introduced in Junos OS Release 11.2.</p>
<b>Description</b>	<p>For MX Series routers with MPCs and T4000 routers with Type 5 FPCs , select data used in the hash key for enhanced IP forwarding engines.</p>

By default, MPCs use the following parameters for hashing:

- Source IP address
- Destination IP address
- Layer 3 protocol
- Source port
- Destination port
- Generic routing encapsulation (GRE) for GRE packets only.

You can modify the default hashing mechanism on MPCs and Type 5 FPCs by configuring statements at the **[edit forwarding-options enhanced-hash-key]** hierarchy level.

**Default** Not enabled.

**Options** **services-loadbalancing**—Distributes traffic across PICs based on source IP address when a route pointing to more than one services PICs is installed.

Data selections for **services-loadbalancing**:

- **inet**—IPv4 addressing protocol.
- **inet6**—IPv6 addressing protocol.
- **layer-3-services**—Include layer 3 IP data in the hash key.
- **incoming-interface-index**—Include incoming interface index in the hash key.
- **source-address**—Include source-address in the hash key.
- **destination-address**—Include destination-address in the hash key.

Data selections for family **inet**:

- **incoming-interface-index**—Include incoming interface index in the hash key.
- **no-destination-port**—Omit IP destination port in the hash key.
- **no-source-port**—Omit IP source port in the hash key..
- **type-of-service**—Include type-of-service (TOS) byte in the hash key.

Data selections for family **inet6**:

- **incoming-interface-index**—Include the incoming interface index in the hash key.
- **no-destination-port**—Omit the IP destination port in the hash key.
- **no-source-port**—Omit the IP source port in the hash key.
- **traffic-class**—Include the traffic class byte in the hash key.

Data selections for family **mpls**:

- **incoming-interface-index**—Include the incoming interface index in the hash key.
- **label-1-exp**—The EXP bit of the first label is used in the hash calculation.
- **no-payload**—Omit the MPLS payload data from the hash key.

Data selections for family **multiservice**:

- **incoming-interface-index**—Include the incoming interface index in the hash key.
- **no-payload**—Omit the payload data from the hash key.
- **outer-priority**—Include the outer 802.1 priority bits in the hash key.

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

**Related Documentation**

- [edit forwarding-options] Hierarchy Level

## family (Filtering)

**Syntax**    `family family-name {  
                   filter {  
                     input input-filter-name;  
                     output output-filter-name;  
                   }  
                   flood {  
                     input filter-name;  
                   }  
                   route-accounting;  
                   source-checking;  
                 }`

**Hierarchy Level**    [edit [forwarding-options](#)]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                               **route-accounting** option introduced in Junos OS Release 8.3; supported only with IPv6.  
                               **source-checking** option introduced in Junos OS Release 12.3 on MX Series Universal Edge Routers; 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 MPLS, or **vpls** for virtual private LAN service (VPLS).



**NOTE:** In Junos OS Release 8.4 and later, the **output** statement is not valid at the [edit [forwarding-options family vpls filter](#)] hierarchy level.

The remaining statements are explained separately.

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

**Related Documentation**    • [Applying Filters to Forwarding Tables on page 231](#)

## family (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 OS Release 7.4.

**Description** Configure flow monitoring for an address family. Only the IPv4 protocol is supported.  
The remaining statements are explained separately.

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

**Related Documentation**

- [Configuring Flow Monitoring on page 235](#)

## family (Port Mirroring)

---

<b>Syntax</b>	<pre>family (ccc   inet   inet6   vpls) {     output {         interface <i>interface-name</i> {             next-hop <i>address</i>;         }         no-filter-check;     } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options port-mirroring</a> ], [edit <a href="#">forwarding-options port-mirroring instance</a> <i>instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>vpls</b> option introduced in Junos OS Release 9.3 for MX Series routers only; support extended to M7i, M10i, M120, and M320 routers in Junos OS Release 9.5. <b>ccc</b> option introduced in Junos OS Release 9.6 for M120 and M320 routers only.
<b>Description</b>	Configure the address type family to sample for port mirroring.
<b>Options</b>	<b>ccc</b> —Sample Layer 2 VPN traffic.  <b>inet</b> —Sample IPv4 traffic.  <b>inet6</b> —Sample IPv6 traffic.  <b>vpls</b> —Sample VPLS traffic  The remaining statements are explained separately.
<b>Required Privilege Level</b>	<b>interface</b> —To view this statement in the configuration. <b>interface-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Port Mirroring on page 246</a></li></ul>

## family (Sampling)

```
Syntax  family (inet | inet6 | mpls) {
        disable;
        output {
            aggregate-export-interval seconds;
            extension-service service-name;
            file {
                disable;
                filename filename;
                files number;
                size bytes;
                (stamp | no-stamp);
                (world-readable | no-world-readable);
            }
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            flow-server 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;
                version9 {
                    template template-name;
                }
            }
            interface interface-name {
                engine-id number;
                engine-type number;
                source-address address;
            }
        }
    }
```

**Hierarchy Level** [edit [forwarding-options sampling](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
**mpls** option introduced in Junos OS Release 8.3.

**Description** Configure the protocol family to be sampled.

**Options** **inet**—IP version 4 (IPv4)  
**inet6**—IP version 6 (IPv6)

## **mpls—MPLS**

The remaining statements are explained separately.

<b>Required Privilege</b>	interface—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Traffic Sampling on page 213</a></li> </ul>



## family inet

<b>Syntax</b>	<pre>family inet {   layer-3;   layer-4;   session-id;   symmetric-hash {     complement;   } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options hash-key</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure layer information for the load-balancing specification. Only the IPv4 protocol is supported.
<b>Options</b>	<p><b>family inet</b>—Incorporate port data into the hash key for flow determination. By default, port data is ignored when determining flows.</p> <ul style="list-style-type: none"> <li>• <b>layer-3</b>—Incorporate Layer 3 (IP) data into the hash key. You must include the <b>layer-3</b> statement. If you omit the <b>layer-3</b> statement, the management process removes the <b>hash-key</b> statement from the configuration and the router behaves as if you specified <b>layer-3</b>.</li> </ul> <p>By default, or if you specify only the <b>layer-3</b> statement, the router uses the following Layer 3 information in the packet header for per-flow load balancing:</p> <ul style="list-style-type: none"> <li>• Source IP address</li> <li>• Destination IP address</li> <li>• Protocol</li> <li>• Incoming interface index</li> </ul> <ul style="list-style-type: none"> <li>• <b>layer-4</b>—Incorporate Layer 4 Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) data into the hash key.</li> </ul> <p>If you include the <b>layer-4</b> statement, the router uses the following Layer 4 information to load-balance:</p> <ul style="list-style-type: none"> <li>• Source port number</li> <li>• Destination port number</li> <li>• IP type of service</li> </ul> <ul style="list-style-type: none"> <li>• <b>session-id</b>—Include the session ID in the hash key.</li> <li>• <b>symmetric-hash</b>—Create the symmetric hash key with source and destination ports. <ul style="list-style-type: none"> <li>• <b>complement</b>—Create the complement of the symmetric hash key.</li> </ul> </li> </ul>

<b>Required Privilege</b>	interface—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Per-Packet Load Balancing on page 186</a></li></ul>

## family mpls

**Syntax**

```
family mpls {
  all-labels;
  label-1;
  label-2;
  label-3;
  no-labels;
  no-label-1-exp;
  payload {
    ether-pseudowire;
    ip {
      disable;
      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 OS Release 7.4.  
**no-label-1-exp** option introduced in Junos OS Release 8.0.  
**label-3** and **no-labels** options introduced in Junos OS Release 8.1.  
**ether-pseudowire** option introduced in Junos OS Release 9.1 (M320 and T Series routers only); support extended to M120 and MX Series routers in Junos OS Release 9.4.  
**all-labels** and **payload ip disable** options introduced in Junos OS Release 12.1X48R2. (PTX Series Packet Transport Switches only).

**Description** For aggregated Ethernet and SONET/SDH interfaces only, configure load balancing based on MPLS labels and payload. Only the IPv4 protocol is supported.

**Options** **family mpls**—(Aggregated Ethernet interfaces, aggregated SONET/SDH interfaces, and multiple equal-cost MPLS next hops only) Incorporate MPLS label and payload information into the hash key for per-flow load balancing. Only the IPv4 protocol is supported.

- **all-labels**—(PTX Series Packet Transport Switches only) Up to eight MPLS labels are included in the hash key to identify the uniqueness of a flow in the Packet Forwarding Engine. This is the default setting.
- **label-1**—(M120, M320, MX Series, and T Series routers only) Include the first MPLS label into the hash key. This is used for a one-label packet for per-flow load balancing IPv4 VPLS traffic based on IP information and MPLS labels.
- **label-2**—(M120, M320, MX Series, and T Series routers only) Include the second MPLS label into the hash key. This is used for a two-label packet for per-flow load balancing

IPv4 VPLS traffic based on IP information and MPLS labels. To use the second MPLS label in the hash key, include both the **label-1** and **label-2** statements at the **[edit forwarding-options hash-key family mpls]** hierarchy level. By default, the router provides hashing on the first and second labels. If both labels are specified, the entire first label and the first 16 bits of the second label are hashed.

- **label-3**—(M120, M320, MX Series, and T Series routers only) Include the third MPLS label into the hash key. To use the third MPLS label, include the **label-1**, **label-2**, and **label-3** statements at the **[edit forwarding-options hash-key family mpls]** hierarchy level.
- **no-labels**—Include no MPLS labels into the hash key.
- **no-label-1-exp**—(M120, M320, MX Series, and T Series routers only) The EXP bit of the first label is not used in the hash calculation to avoid reordering complications.
- **payload**—Incorporate bits from the IP payload into the hash key for per-flow load balancing Layer 2 information based on MPLS labels.
  - **disable**—(PTX Series Packet Transport Switches only) Exclude IP payload from the hash key.
  - **ether-pseudowire**—(M120, M320, MX Series, and T Series routers only) Load-balance IPv4 traffic over Layer 2 Ethernet pseudowires.
  - **ip**—Include the IP address of the IPv4 or IPv6 payload into the hash key for per-flow load balancing Layer 2 information based on MPLS labels. For the PTX Series Packet Transport Switches, this is the default setting with both Layer 3 and Layer 4 IP information included in the hash key.
    - **disable**—(PTX Series Packet Transport Switches only) Exclude IP payload from the hash key.
    - **layer-3-only**—Include only Layer 3 IP information from the IP payload data into the hash key for per-flow load balancing Layer 2 information based on MPLS labels.
    - **port-data**—(M120, M320, MX Series, and T Series routers only) Include the source and destination port field information into the hash key. By default, the most significant byte and least significant byte of the source and destination port fields are hashed. To select specific bytes 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 bytes from being hashed, include the **layer-3-only** statement at the **[edit forwarding-options hash-key family mpls payload ip]** hierarchy level.
      - **destination-lsb**—Include the least-significant byte of the destination port.
      - **destination-msb**—Include the most-significant byte of the destination port.
      - **source-lsb**—Include the least-significant byte of the source port.
      - **source-msb**—Include the most-significant byte of the source port.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
---------------------------------	---

- Related Documentation**
- [Configuring Load Balancing Based on MPLS Labels on DPC I-Chip-Based Hardware on page 189](#)
  - [Configuring Load Balancing for Ethernet Pseudowires on page 193](#)

## family multiservice

```
Syntax  family multiservice {
        destination-mac;
        label-1;
        label-2;
        payload {
            ip {
                layer-3 {
                    (source-ip-only | destination-ip-only);
                }
                layer-3-only;
                layer-4;
            }
        }
        source-mac;
        symmetric-hash {
            complement;
        }
    }
```

**Hierarchy Level** [edit [forwarding-options hash-key](#)]

**Release Information** Statement introduced in Junos OS Release 8.0.  
**ip**, **label-1**, **label-2**, **layer-3-only**, and **payload** options introduced in Junos OS Release 9.4.  
**layer-3**, **layer-4**, **source-ip-only**, and **destination-ip-only** options introduced in Junos OS Release 9.5.  
**symmetric-hash** and **complement** options introduced in Junos OS Release 9.6.

**Description** Configure load balancing based on Layer 2 media access control information. On MX Series routers, configure VPLS load balancing. On M120 and M320 routers only, configure VPLS load balancing based on MPLS labels and IP information. For IPv4 traffic, only the IP source and destination addresses are included in the hash key. For MPLS and IPv4 traffic, one or two MPLS labels and IPv4 source and destination addresses are included. For MPLS Ethernet pseudowires, only one or two MPLS labels are included in the hash key.

**Options** You can configure one or more options to load-balance using the packet information that you specify.

**destination-mac**—Include the destination-address MAC information in the hash key for Layer 2 load balancing.

**label-1** (M120 and M320 routers only)—Include the first MPLS label in the hash key. Used for including a one-label packet for per-flow load balancing of IPv4 VPLS traffic based on IP information and MPLS labels.

**label-2** (M120 and M320 routers only)—Include the second MPLS label in the hash key. If both **label-1** and **label-2** are specified, the entire first label and the first 16 bits of the second label are hashed.

**payload** (MX Series, M120, and M320 routers only)—Include the packet's IP payload in the hash key.

- **ip** (MX Series, M120, and M320 routers only)—Include the IP address of the IPv4 or IPv6 payload in the hash key.
- **layer-3** (MX Series routers only)—Use this to include Layer 3 information from the packet's IP payload in the hash key.
  - **destination-ip-only** (MX Series routers only)—Use this to include only the destination IP address in the payload in the hash key.
  - **source-ip-only** (MX Series routers only)—Use this to include only the source IP address in the payload in the hash key.



**NOTE:** You can include either the **source-ip-only** or the **destination-ip-only** statement, not both. They are mutually exclusive.

- **layer-3-only** (M120, and M320 routers only)—Include only the Layer 3 information from the packet's IP payload in the hash key.
- **layer-4** (MX Series routers only)—Include Layer 4 information from the packet's IP payload in the hash key.



**NOTE:** On MX Series routers only, you can configure either Layer 3 or Layer 4 load balancing, or both at the same time.



**NOTE:** On I chip platforms, an unknown Layer 4 header is excluded from load-balance hashing to avoid undesired packet reordering.

**source-mac**—Include the source-address MAC information in the hash key.

**symmetric-hash** (MX Series routers only)—Configure the symmetric hash or symmetric hash complement for configuring symmetrical load balancing on an 802.3ad Link Aggregation Group.

- **complement** —Include the complement of the symmetric hash in the hash key.

**Required Privilege Level**

interface—To view this statement in the configuration.

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

- Related Documentation**
- [Configuring Load Balancing Based on MAC Addresses on page 194](#)
  - [Configuring VPLS Load Balancing Based on IP and MPLS Information on page 195](#)
  - [Configuring VPLS Load Balancing on MX Series 3D Universal Edge Routers on page 196](#)
  - [Configuring VPLS Load Balancing](#)

---

## file (Extended DHCP Relay Agent and Helpers Trace Options)

---

<b>Syntax</b>	file <i>filename</i> <files <i>number</i> > <match <i>regular-expression</i> > <size <i>bytes</i> > <world-readable   no-world-readable>;
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options dhcp-relay traceoptions</a> ], [edit <a href="#">forwarding-options helpers traceoptions</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure information about the DNS and TFTP packet-forwarding files that contain trace logging information.
<b>Options</b>	<i>filename</i> —Name of the file containing the trace information. <b>Default:</b> /var/log/sampled  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <a href="#">Tracing BOOTP, DNS, and TFTP Forwarding Operations on page 242</a>

---

## file (Packet Capture)

---

<b>Syntax</b>	file <i>filename filename</i> <files <i>number</i> > <size <i>bytes</i> > <world-readable   no-world-readable>;
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options packet-capture</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Enable packet capture to a file.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <a href="#">Configuring Packet Capture on page 250</a>



## file (Sampling)

---

<b>Syntax</b>	file <i>filename filename</i> <disable> <files <i>number</i> > <stamp   no-stamp> <size <i>bytes</i> > <world-readable   no-world-readable>;
<b>Hierarchy Level</b>	[edit forwarding-options sampling family <i>family-name</i> output]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Collect the traffic samples in a file.  The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Output File for Traffic Sampling on page 215</a></li> </ul>

## file (Trace Options)

---

<b>Syntax</b>	file <i>filename</i> <files <i>number</i> > <size <i>bytes</i> > <world-readable   no-world-readable>;
<b>Hierarchy Level</b>	[edit forwarding-options port-mirroring traceoptions], [edit forwarding-options sampling traceoptions]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure information about the files that contain trace logging information.
<b>Options</b>	<i>filename</i> —The name of the file containing the trace information. <b>Default:</b> /var/log/sampled  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Tracing Traffic-Sampling Operations on page 217</a></li> </ul>

## filename (Packet Capture)

---

<b>Syntax</b>	<code>filename <i>filename</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options packet-capture file</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Configure the name of the output file.
<b>Options</b>	<i>filename</i> —Name of the file.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Packet Capture on page 250</a></li></ul>

## filename (Sampling)

---

<b>Syntax</b>	<code>filename <i>filename</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options sampling family <i>family-name</i> output file</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the name of the output file.
<b>Options</b>	<i>filename</i> —Name of the file in which to place the traffic samples. All files are placed in the directory <code>/var/tmp</code> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Output File for Traffic Sampling on page 215</a></li></ul>

## files (Packet Capture)

---

<b>Syntax</b>	<code>files <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options packet-capture file</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Configure the maximum number of files for packet capturing.
<b>Options</b>	<i>number</i> —Maximum number of files. <b>Range:</b> 2 through 10,000 files <b>Default:</b> 10
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Packet Capture on page 250</a></li></ul>

## files (Sampling and Traceoptions)

---

<b>Syntax</b>	<code>files <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpertraceoptions file</a> ], [edit <a href="#">forwarding-options port-mirroring traceoptions file</a> ], [edit <a href="#">forwarding-options sampling family <i>family-name</i> output file</a> ], [edit <a href="#">forwarding-options sampling traceoptions file</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.
<b>Description</b>	Configure the total number of files to be saved with samples or trace data.
<b>Options</b>	<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. <b>Range:</b> 1 through 100 files <b>Default:</b> 5 files for sampling output; 10 files for trace log information
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Output File for Traffic Sampling on page 215</a></li><li>• <a href="#">Tracing Traffic-Sampling Operations on page 217</a></li></ul>

## filter (IPv4, IPv6, and MPLS)

---

<b>Syntax</b>	<pre>filter {     input input-filter-name;     output output-filter-name; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options family</a> (inet   inet6   mpls)]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply a forwarding table filter to a forwarding table.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Applying Filters to Forwarding Tables on page 231</a></li></ul>

## filter (VPLS)

---

<b>Syntax</b>	<pre>filter input filter-name;</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options family</a> vpls], [edit logical-systems <i>logical-system-name</i> <a href="#">forwarding-optionsfamily</a> vpls], [edit logical-systems <i>logical-system-name</i> routing -instances <i>instance-name</i> <a href="#">forwarding-optionsfamily</a> vpls], [edit routing -instances <i>instance-name</i> <a href="#">forwarding-optionsfamily</a> vpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply a forwarding table filter for VPLS.
<b>Options</b>	The other statement is explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Applying Filters to Forwarding Tables on page 231</a></li></ul>

## flood

---

<b>Syntax</b>	<code>flood {     input <i>filter-name</i>; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-optionsfamily vpls</a> ], [edit routing -instances <i>instance-name</i> <a href="#">forwarding-optionsfamily vpls</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply a forwarding table filter to a flood table.
<b>Options</b>	<code>input <i>filter-name</i></code> —Name of the forwarding table filter.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Applying Filters to Forwarding Tables on page 231</a></li></ul>

## flow-active-timeout

---

<b>Syntax</b>	<code>flow-active-timeout <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options accounting group-name output</a> ], [edit <a href="#">forwarding-options monitoring group-name family inet output</a> ], [edit <a href="#">forwarding-optionssampling family family-name output</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the time that elapses before another active flow is exported.
<b>Options</b>	<code><i>seconds</i></code> —Timeout, in seconds. <b>Range:</b> 60 through 1800 <b>Default:</b> 1800
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Discard Accounting on page 234</a></li><li>• <a href="#">Configuring Flow Monitoring on page 235</a></li><li>• <a href="#">Configuring the Output File for Traffic Sampling on page 215</a></li></ul>

## flow-export-destination

---

<b>Syntax</b>	<code>flow-export-destination {     (cflowd-collector   collector-pic); }</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options monitoring group-name family inet output</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure flow collection.
<b>Options</b>	<code>cflowd-collector</code> —cflowd collector.  <code>collector-pic</code> —Collector PIC.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flow Monitoring on page 235</a></li></ul>

## flow-inactive-timeout

---

<b>Syntax</b>	<code>flow-inactive-timeout <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options accounting group-name output</a> ], [edit <a href="#">forwarding-options monitoring group-name family inet output</a> ], [edit <a href="#">forwarding-options sampling family family-name output</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the time that elapses before a flow is considered inactive.
<b>Options</b>	<code>seconds</code> —Timeout, in seconds. <b>Range:</b> 15 through 1800 <b>Default:</b> 60
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Discard Accounting on page 234</a></li><li>• <a href="#">Configuring Flow Monitoring on page 235</a></li><li>• <a href="#">Configuring the Output File for Traffic Sampling on page 215</a></li></ul>

## flow-server

<b>Syntax</b>	<pre> flow-server <i>hostname</i> {     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 <i>port-number</i>;     source-address <i>address</i>;     version <i>format</i>;     version9 {         template <i>template-name</i>;     } } </pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options sampling family <i>family-name</i> output</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>version9</b> statement introduced in Junos OS Release 8.3.</p>
<b>Description</b>	<p>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.</p> <p>You can configure up to one version 5 and one version 8 flow format at the [edit <a href="#">forwarding-options sampling output flow-server <i>hostname</i></a>] 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.</p>
<b>Options</b>	<p><b><i>hostname</i></b>—The IP address or identifier of the host system (the workstation either running the cflowd utility or collecting traffic flows using version 9).</p> <p>You can configure only one host system for version 9.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Active Flow Monitoring Using Version 9 on page 220</a></li> <li>• <a href="#">Configuring Flow Aggregation (cflowd) on page 217</a></li> </ul>

## forwarding-options

---

<b>Syntax</b>	<code>forwarding-options { ... }</code>
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Configure traffic forwarding.  The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Traffic Forwarding and Monitoring on page 225</a></li><li>• [edit forwarding-options] Hierarchy Level</li></ul>

## group (DHCP Spoofing Prevention)

---

<b>Syntax</b>	<code>group <i>group-name</i> {     interface <i>interface-name</i>; }</code>
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4 (MX Series routers only).
<b>Description</b>	Configure Dynamic Host Configuration Protocol (DHCP) snooping on the router. When acting as a snooping agent, the MX Series router typically is located between the client and the DHCP relay agent. It creates filters by “snooping” DHCP messages and binding DHCP-issued IP addresses with the MAC address of the client. These filters help prevent DHCP spoofing.  Configure DHCP snooping by including the appropriate interfaces under the <b>group</b> statement.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Preventing DHCP Spoofing on MX Series 3D Universal Edge Routers on page 245</a></li></ul>



## hash-key (Forwarding Options)

```
Syntax hash-key {
    family inet {
        layer-3;
        layer-4;
        session-id;
        symmetric-hash {
            complement;
        }
    }
    family mpls {
        all-labels;
        label-1;
        label-2;
        label-3;
        no-labels;
        no-label-1-exp;
        payload {
            ether-pseudowire;
            ip {
                disable;
                layer-3-only;
                port-data {
                    destination-lsb;
                    destination-msb;
                    source-lsb;
                    source-msb;
                }
            }
        }
    }
}
family multiservice {
    destination-mac;
    label-1;
    label-2;
    payload {
        ip {
            layer-3-only;
            layer-3 {
                (source-ip-only | destination-ip-only);
            }
            layer-4;
        }
    }
    source-mac;
}
```

Hierarchy Level [edit [forwarding-options](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
**family multiservice** and **no-label-1-exp** options introduced in Junos OS Release 8.0.  
**label-3** and **no-labels** options introduced in Junos OS Release 8.1.

**ether-pseudowire** statement introduced in Junos OS Release 9.1 (M320 and T Series routers only); support extended to M120 and MX Series routers in Junos OS Release 9.4. **ip**, **label-1**, **label-2**, **layer-3-only**, and **payload** options for the **family multiservice** statement introduced in Junos OS Release 9.4. (M120 and M320 routers only); support for **ip** and **payload** statements only to MX Series routers.

**layer-3**, **source-ip-only**, **destination-ip-only**, and **layer-4** statements introduced for the **family multiservice** statement in Junos OS Release 9.5. (MX Series routers only).

**all-labels** and **payload ip disable** statements introduced for the **family mpls** statement in Junos OS Release 12.1X48R2. (PTX Series Packet Transport Switches only).

**Description** Select which packet header data to use for per-flow load balancing.

The options are explained separately.



**NOTE:** To modify the default hashing mechanism on Modular Port Concentrators (MPCs) and Type 5 FPCs, you need to configure the statements at the [edit forwarding-options enhanced-hash-key] hierarchy level. Statements at the [edit forwarding-options hash-key] hierarchy level do not support MPCs and Type 5 FPCs.



**NOTE:**

The following statements are not supported on T Series routers:

- The **symmetric-hash** and the **session-id** statements at the [edit forwarding-options hash-key family inet] hierarchy level and all statements at the [edit forwarding-options hash-key family multiservice] hierarchy level.
- The **label-1** and **label-2** statements, and the IP configuration at the [edit forwarding-options hash-key family multiservice] hierarchy level.

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

**Related Documentation**

- [Configuring Per-Packet Load Balancing on page 186](#)
- [Configuring Load Balancing Based on MPLS Labels on DPC I-Chip-Based Hardware on page 189](#)
- [Configuring Load Balancing Based on MAC Addresses on page 194](#)

## helpers

```
Syntax  helpers {
    bootp {
        client-response-ttl number;
        description text-description;
        interface interface-group {
            client-response-ttl number;
            description text-description;
            maximum-hop-count number;
            minimum-wait-time seconds;
            no-listen;
            server address {
                logical-system logical-system-name <routing-instance [ <default>
                    routing-instance-names ]>;
                routing-instance [ <default> routing-instance-names ];
            }
        }
        maximum-hop-count number;
        minimum-wait-time seconds;
        relay-agent-option;
        server address {
            logical-system logical-system-name <routing-instance [ <default>
                routing-instance-names ]>;
            routing-instance [ <default> routing-instance-names ];
        }
    }
    domain {
        description text-description;
        interface interface-name {
            broadcast;
            description text-description;
            no-listen;
            server address <logical-system logical-system-name> <routing-instance
                routing-instance-name>;
        }
        server address <logical-system logical-system-name> <routing-instance
            routing-instance-name>;
    }
    port (Packet Forwarding) port-number {
        description text-description;
        interface interface-name {
            broadcast;
            description text-description;
            no-listen;
            server address <logical-system logical-system-name> <routing-instance
                routing-instance-name>;
        }
        server address <logical-system logical-system-name> <routing-instance
            routing-instance-name>;
    }
    tftp {
        description text-description;
        interface interface-name {
```

```

    broadcast;
    description text-description;
    no-listen;
    server address <logical-system logical-system-name> <routing-instance
        routing-instance-name>;
  }
  server address <logical-system logical-system-name> <routing-instance
    routing-instance-name>;
}
traceoptions {
  file filename <files number> <match regular-expression> <size bytes> <world-readable |
    no-world-readable>;
  flag flag;
  level level;
  no-remote-trace level;
}
}

```

**Hierarchy Level** [edit [forwarding-options](#)]

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

**Description** Enable TFTP or DNS request packet forwarding, or configure the router, switch, or interface to act as a DHCP/BOOTP relay agent. Use only one server address per interface or global configuration.

In addition to TFTP and DNS, you can use the helpers statement to enable LAN-Broadcast forwarding. To do this, you must manually define which UDP port is forwarded, as shown here:

```

user@ host# show forwarding-options
helpers {
  port 3000 {
    interface {
      fe-0/0/1.0 {
        server 111.0.0.2;
      }
    }
  }
  port 3001 {
    interface {
      fe-0/0/0.0 {
        server 100.0.0.2;
      }
    }
  }
}


```

The remaining statements are explained separately.

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

- Related Documentation**
- [Configuring DNS and TFTP Packet Forwarding on page 241](#)
  - [Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239](#)

## indexed-load-balance

<b>Syntax</b>	indexed-load-balance;
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options load-balance</a> ], [edit logical-systems <i>logical-system-name</i> <a href="#">forwarding-options load-balance</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <a href="#">forwarding-options load-balance</a> ], [edit routing-instances <i>routing-instance-name</i> <a href="#">forwarding-options load-balance</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Improve load-balance distribution for unicast and aggregated next hops. Include this statement if you notice issues with load-balance distribution for IPv4 traffic. The <b>indexed-load-balance</b> statement causes the creation of a nexthop structure that is not a function of the hash only, but is also a function of the low-order bits of the IP address.
<div>  <p><b>CAUTION:</b> Including the <b>indexed-load-balance</b> statement causes an increase in memory usage on the device.</p> </div>	
<b>Default</b>	Disabled
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Per-Prefix Load Balancing on page 238</a></li> </ul>

## input (Forwarding Table)

---

<b>Syntax</b>	<code>input <i>filter-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options family</a> (inet   inet6   mpls   vpls) <a href="#">filter</a> ], [edit routing-instances <i>routing-instance-name</i> <a href="#">forwarding-options family</a> (inet   inet6   mpls   vpls) <a href="#">filter</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply a forwarding table filter to ingress traffic of the forwarding table.
<b>Options</b>	<i>filter-name</i> —Name of the applied filter.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Applying Filters to Forwarding Tables on page 231</a></li></ul>

## input (Port Mirroring)

---

<b>Syntax</b>	<code>input {     <a href="#">maximum-packet-length</a> <i>bytes</i>;     <a href="#">rate</a> <i>number</i>;     <a href="#">run-length</a> <i>number</i>; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options port-mirroring</a> ], [edit <a href="#">forwarding-options port-mirroring instance</a> <i>instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>maximum-packet-length</b> option introduced in Junos OS Release 9.6 for M120 and M320 routers only. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.
<b>Description</b>	Configure input packet properties for port mirroring.  The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Port Mirroring on page 246</a></li></ul>

## input (Sampling)

---

<b>Syntax</b>	<pre>input {   max-packets-per-second <i>number</i>;   maximum-packet-length <i>bytes</i>;   rate <i>number</i>;   run-length <i>number</i>; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options sampling</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Support for sampling of MPLS traffic introduced in Junos OS Release 8.3.
<b>Description</b>	Configure traffic sampling on a logical interface.  The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Traffic Sampling on page 213</a></li></ul>

## instance

---

**Syntax**

```
instance {  
  instance-name {  
    input {  
      maximum-packet-length bytes;  
      rate number;  
      run-length number;  
    }  
    family (ccc| inet | inet6 | vpls) {  
      output {  
        interface interface-name {  
          next-hop address;  
        }  
        no-filter-check;  
      }  
    }  
  }  
}
```

**Hierarchy Level** [edit [forwarding-options port-mirroring](#)]

**Release Information** Statement introduced in Junos OS Release 9.3 (MX Series routers only). Support extended to M120 and M320 routers in Junos OS Release 9.5. **maximum-packet-length** and **ccc** options introduced in Junos OS Release 9.6 for M120 and M320 routers only.

**Description** Configure a port-mirroring instance.

**Options** *port-mirroring-instance-name*—Name of the port-mirroring instance.

The remaining statements are explained separately.

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

**Related Documentation**

- [Configuring Port Mirroring on page 246](#)



## interface (Accounting or Sampling)

---

<b>Syntax</b>	<pre>interface <i>interface-name</i> {     engine-id <i>number</i>;     engine-type <i>number</i>;     source-address <i>address</i>; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options accounting group-name output</a> ], [edit <a href="#">forwarding-options sampling family family-name output</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the output interface for sending copies of packets elsewhere to be analyzed.
<b>Options</b>	<p><b>engine-id <i>number</i></b>—Identity of the accounting interface.</p> <p><b>engine-type <i>number</i></b>—Type of this accounting interface.</p> <p><b><i>interface-name</i></b>—Name of the accounting interface.</p> <p><b>source-address <i>address</i></b>—Address used for generating packets.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Discard Accounting on page 234</a></li><li>• <a href="#">Configuring the Output File for Traffic Sampling on page 215</a></li></ul>

## interface (BOOTP)

---

<b>Syntax</b>	<pre>interface (<i>interface-name</i>   <i>interface-group</i>) {     broadcast;     client-response-ttl <i>number</i>;     description <i>text-description</i>;     maximum-hop-count <i>number</i>;     minimum-wait-time <i>seconds</i>;     no-listen;     server <i>address</i> {         logical-system <i>logical-system-name</i> &lt;routing-instance [ &lt;default&gt;             <i>routing-instance-names</i> ]&gt;;         routing-instance [ &lt;default&gt; <i>routing-instance-names</i> ];     }     apply-secondary-as-giaddr (QFX platforms only) }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpers bootp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Specify the interface for a DHCP and BOOTP relay agent.
<b>Options</b>	<p><b><i>interface-group</i></b>—Sets a logical interface or group of logical interfaces with a specific DHCP relay configuration.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239</a></li><li>• <a href="#">Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)</a></li></ul>

## interface (DHCP Spoofing Prevention)

---

<b>Syntax</b>	<code>interface <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-relay group <i>group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4 (MX Series routers only).
<b>Description</b>	<p>Configure Dynamic Host Configuration Protocol (DHCP) snooping on the router. When acting as a snooping agent, the MX Series router typically is located between the client and the DHCP relay agent. It creates filters by “snooping” DHCP messages and binding DHCP-issued IP addresses with the MAC address of the client. These filters help prevent DHCP spoofing.</p> <p>DHCP snooping is configured by including the appropriate interfaces.</p>
<b>Required Privilege Level</b>	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Preventing DHCP Spoofing on MX Series 3D Universal Edge Routers on page 245</a></li></ul>

## interface (DNS and TFTP Packet Forwarding or Relay Agent)

---

<b>Syntax</b>	<pre>interface <i>interface-name</i> {     broadcast;     <b>description</b> <i>text-description</i>;     <b>no-listen</b>;     <b>server</b> <i>address</i> &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance         <i>routing-instance-name</i>&gt;; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpers domain</a> ], [edit <a href="#">forwarding-options helpers tftp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the interface for monitoring and forwarding DNS or TFTP requests.
<b>Options</b>	<p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring DNS and TFTP Packet Forwarding on page 241</a></li></ul>

## interface (Monitoring)

---

<b>Syntax</b>	<pre>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>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options monitoring group-name</a> inet <a href="#">output</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the output interface for monitored traffic.
<b>Options</b>	<p><i>interface-name</i>—Name of the interface.</p> <p><i>engine-id number</i>—Identity of the monitoring interface.</p> <p><i>engine-type number</i>—Type of this monitoring interface.</p> <p><i>input-interface-index number</i>—Input interface index for records from this interface.</p> <p><i>output-interface-index number</i>—Output interface index for records from this interface.</p> <p><i>source-address address</i>—Address used for generating packets.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flow Monitoring on page 235</a></li></ul>

## interface (Next-Hop Group)

---

<b>Syntax</b>	<code>interface <i>interface-name</i> {     <b>next-hop</b> <i>address</i>; }</code>
<b>Hierarchy Level</b>	[edit <b>forwarding-options next-hop-group</b> <i>group-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>Specify the output interface for sending copies of packets elsewhere to be analyzed.</p> <p>The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.</p>
<b>Options</b>	<p><b><i>interface-name</i></b>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Next-Hop Groups on page 237</a></li></ul>

## interface (Port Mirroring)

---

<b>Syntax</b>	<code>interface <i>interface-name</i> {     <b>next-hop</b> <i>address</i>; }</code>
<b>Hierarchy Level</b>	[edit <b>forwarding-options port-mirroring output</b> ], [edit <b>forwarding-options port-mirroring family</b> (inet   inet6) <b>output</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the output interface for sending copies of packets elsewhere to be analyzed.
<b>Options</b>	<p><b><i>interface-name</i></b>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Port Mirroring on page 246</a></li></ul>

## load-balance (Forwarding Options)

<b>Syntax</b>	<pre>load-balance {   indexed-load-balance;   per-flow {     hash-seed;   }   per-prefix {     hash-seed <i>number</i>;   } }</pre>
<b>Hierarchy Level</b>	<pre>[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]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for per-flow load balancing introduced in Junos OS Release 9.3.</p>
<b>Description</b>	<p>Enable per-prefix or per-flow load balancing so that the router or switch elects a next hop independently of the route selected by other routers or switches.</p> <p>For the active route, when there are multiple equal-cost paths to the same destination, by default, the Junos OS chooses in a random fashion one of the next-hop addresses to install into the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is chosen again, also in a random fashion.</p> <p>You can configure the Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This is called per-packet load balancing. You can use load balancing to spread traffic across multiple paths between routing devices. The behavior of the per-packet load-balancing function varies according to the version of the Internet Processor ASIC in the routing device.</p> <p>On routing devices with an Internet Processor I ASIC, when per-packet load balancing is configured, traffic between routing devices with multiple paths is spread in a random fashion across the available interfaces. The forwarding table balances the traffic headed to a destination, transmitting packets 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.</p> <p>Per-packet load distribution uses a hashing algorithm that distributes packets over equal-cost links. The algorithm is designed to distribute packets to prevent any single link from being saturated. However, per-packet load balancing offers no guarantee of equal distribution of traffic over equal-cost links, nor does it guarantee that increasing the number of Internet flows creates a better hash distribution.</p> <p>On routing devices with the Internet Processor II ASIC and T Series Internet Processor II ASIC, when per-packet load balancing is configured, traffic between routing devices with multiple paths is divided into individual traffic flows (up to a maximum of 16 equal-cost</p>

load-balanced paths). On some platforms, you can increase the number of paths by using the **chassis maximum-ecmp** statement. With this statement, you can change the maximum number of equal-cost load-balanced paths to 32 or 64. Packets for each individual flow are kept on a single interface. To recognize individual flows in the transit traffic, the routing device examines each of the following:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Source interface index
- Type of service (ToS)

The routing device recognizes packets in which all of these parameters are identical, and it 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.

Load balancing is not supported on management and internal Ethernet (**fxo**) interfaces because this type of interface cannot handle the routing process. On **fxp** interfaces, you cannot configure multiple next hops and enable load balancing.

**Options**     The statements are explained separately.

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

**Related Documentation**

- Example: Load Balancing BGP Traffic
- [Configuring Per-Flow Load Balancing Based on Hash Values on page 239](#)
- [Configuring Per-Prefix Load Balancing on page 238](#)



## load-balance-group

---

<b>Syntax</b>	<code>load-balance-group <i>group-name</i> {     next-hop-group [ <i>group-names</i> ]; }</code>
<b>Hierarchy Level</b>	[edit firewall]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure a load-balance group.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Load-Balance Groups on page 237</a> in the Routing Policy Configuration Guide</li></ul>

## local-dump

---

<b>Syntax</b>	<code>(local-dump   no-local-dump);</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options sampling family</a> (inet   inet6   mpls) <a href="#">output flow-server</a> <i>hostname</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Enable collection of cflowd records in a log file.
<b>Options</b>	<i>no-local-dump</i> —Do not dump cflowd records to a log file before exporting.  <i>local-dump</i> —Dump cflowd records to a log file before exporting.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Debugging cflowd Flow Aggregation on page 219</a></li></ul>

## max-packets-per-second

---

<b>Syntax</b>	max-packets-per-second <i>number</i> ;
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options sampling input</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<i>number</i> —Maximum number of packets per second. <b>Range:</b> 0 through 65,535 <b>Default:</b> 1000
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• See <a href="#">Configuring Traffic Sampling on page 213</a>.</li></ul>

## maximum-capture-size

---


<b>Syntax</b>	maximum-capture-size <i>bytes</i> ;
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options packet-capture</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Configure the maximum size of capture for packets.
<b>Options</b>	<i>bytes</i> —Maximum capture size. <b>Range:</b> 68 through 1500 <b>Default:</b> 68 bytes
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Packet Capture on page 250</a></li></ul>

## maximum-hop-count

---

<b>Syntax</b>	<code>maximum-hop-count <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpers bootp</a> ], [edit <a href="#">forwarding-options helpers bootpinterface</a> ( <i>interface-name</i>   <i>interface-group</i> )]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Specify the maximum number of hops allowed.
<b>Options</b>	<i>number</i> —Maximum number of hops. <b>Default:</b> 4 hops
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239</a></li></ul>

## maximum-packet-length

<b>Syntax</b>	maximum-packet-length <i>bytes</i> ;
<b>Hierarchy Level</b>	[edit forwarding-options port-mirroring <i>input</i> ], [edit forwarding-options port-mirroring instance <i>instance-name</i> <i>input</i> ], [edit forwarding-options sampling <i>input</i> ], [edit forwarding-options sampling instance <i>instance-name</i> <i>input</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.
<b>Description</b>	Set the maximum length of the packet used for port mirroring or traffic sampling. Packets with lengths greater than the specified maximum are truncated.
<div>  <p><b>NOTE:</b> The maximum-packet-length statement is not supported on MX80 routers.</p> </div>	
<b>Options</b>	<p><i>bytes</i>—Maximum length (in bytes) of the mirrored packet or the sampled packet.</p> <p><b>Range:</b> 0 through 9216</p> <p><b>Default:</b> 0</p> <p>For MX-Series devices with Modular Port Concentrators (MPCs), port-mirrored or sampled packets can be truncated (or clipped) to any length in the range of 1 to 255 bytes. Only 1 to 255 are valid values for packet truncation on these devices. For other devices, the range is from 0 to 9216. A maximum-packet-length value of zero represents that truncation is disabled, and the entire packet is mirrored or sampled.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Port Mirroring on page 246</a></li> <li>• <a href="#">Configuring Traffic Sampling on page 213</a></li> </ul>

## minimum-wait-time

---

<b>Syntax</b>	<code>minimum-wait-time seconds;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpers bootp</a> ], [edit <a href="#">forwarding-options helpers bootpinterface</a> ( <i>interface-name</i>   <i>interface-group</i> )]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Specify the minimum time allowed.
<b>Options</b>	<b>seconds</b> —Minimum time. <b>Default:</b> 0 seconds
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239</a></li></ul>

## mirror-once

---

<b>Syntax</b>	<code>mirror-once;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options port-mirroring</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3 (MX Series routers only). Support extended to M120 routers in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Switches.
<b>Description</b>	Configure the router to mirror packets only once. This feature is useful if you configure port mirroring on both ingress and egress interfaces, which could result in the same packet being mirrored twice.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Port Mirroring on page 246</a></li></ul>

## 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 OS Release 7.4.

**Description** Specify flow monitoring instance name and properties.

The statements are explained separately.

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

**Related Documentation**

- [Configuring Flow Monitoring on page 235](#)

## next-hop (Forwarding Options)

---

<b>Syntax</b>	<code>next-hop <i>address</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options port-mirroring</a> <code>output interface <i>interface-name</i></code> ], [edit <a href="#">forwarding-options port-mirroring family</a> (inet   inet6   ccc   vpls) <code>output interface <i>interface-name</i></code> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.
<b>Description</b>	Specify the next-hop address for sending copies of packets to an analyzer.
<b>Options</b>	<b><i>address</i></b> —IP address of the next-hop router.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Port Mirroring on page 246</a></li></ul>

## next-hop-group

---

**Syntax**    `next-hop-group group-name {  
                  interface interface-name {  
                    next-hop address;  
                  }  
                  next-hop-subgroup subgroup-name {  
                    interface interface-name {  
                      next-hop address;  
                    }  
                  }  
                }`

**Hierarchy Level**    [edit [forwarding-options](#)]

**Release Information**    Statement introduced before Junos OS Release 7.4.

**Description**    Specify the next-hop address for sending copies of packets to an analyzer.

The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.

**Options**    *addresses*—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.

*group-names*—Name of next-hop group. Up to 30 next-hop groups are supported for the router. Each next-hop group must have at least two next-hop addresses.

*interface-name*—Interface used to reach the next-hop destination.

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

**Related Documentation**    • [Configuring Next-Hop Groups on page 237](#)



## no-filter-check

---

<b>Syntax</b>	no-filter-check;
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options port-mirroring output</a> ], [edit <a href="#">forwarding-options port-mirroring family</a> (inet   inet6   ccc   vpls) <a href="#">output</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.
<b>Description</b>	Disable filter checking on the port-mirroring interface.  This statement is required when you send port-mirrored traffic to a Tunnel Services PIC that has a filter applied to it.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Port Mirroring on page 246</a></li> </ul>

## no-listen

---

<b>Syntax</b>	no-listen;
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpers bootp interface</a> ( <i>interface-name</i>   <i>interface-group</i> )], [edit <a href="#">forwarding-options helpers domain interface</a> <i>interface-name</i> ], [edit <a href="#">forwarding-options helpers tftp interface</a> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Disable recognition of DNS requests or stop packets from being forwarded on a logical interface, a group of logical interfaces, a router, or a switch.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring DNS and TFTP Packet Forwarding on page 241</a></li> <li>• <a href="#">Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239</a></li> </ul>

## 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-name {
            engine-id number;
            engine-type number;
            source-address address;
        }
    }
```

**Hierarchy Level** [edit [forwarding-options accounting group-name](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** Configure cflowd, output interfaces, and flow properties.

The statements are explained separately.

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

**Related Documentation**

- [Configuring Discard Accounting on page 234](#)

## output (Forwarding Table)

---

<b>Syntax</b>	<code>output <i>filter-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-optionsfamily</a> (inet   inet6   mpls) <a href="#">filter</a> ], [edit routing-instances <i>routing-instance-name</i> <a href="#">forwarding-optionsfamily</a> (inet   inet6   mpls) <a href="#">filter</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Configure filtering on the egress traffic of the forwarding table.
<b>Options</b>	<i>filter-name</i> —Name of the applied filter.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Applying Filters to Forwarding Tables on page 231</a></li></ul>

## output (Monitoring)

---

**Syntax**

```
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](#)[family inet](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** Configure cflowd, output interfaces, and flow properties.

The statements are explained separately.

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

**Related Documentation**

- [Configuring Flow Monitoring on page 235](#)

## output (Port Mirroring)

---

<b>Syntax</b>	<pre>output {     interface <i>interface-name</i> {         next-hop <i>address</i>;     }     no-filter-check; }</pre>
<b>Hierarchy Level</b>	[edit <b>forwarding-options port-mirroring family</b> (ccc   inet   inet6   vpls)], [edit <b>forwarding-options port-mirroring instance</b> <i>instance-name</i> <b>family</b> (ccc   inet   inet6   vpls)]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>vpls</b> option introduced in Junos OS Release 9.3 for MX Series routers only; support extended to M7i, M10i, M120, and M320 routers in Junos OS Release 9.5. <b>ccc</b> option introduced in Junos OS Release 9.6 for M120 and M320 routers only.
<b>Description</b>	Configure the port mirroring destination properties.  The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Port Mirroring on page 246</a></li></ul>

## output (Sampling)

```
Syntax  output {
        aggregate-export-interval seconds;
        flow-server 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;
            version9 {
                template template-name;
            }
        }
        extension-service service-name;
        file filename filename <disable> <files number> <stamp | no-stamp> <size bytes>
            <world-readable | no-world-readable>;
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        flow-server host-name {
            aggregation;
            autonomous-system-type (origin | peer);
            (local-dump | no-local-dump);
            port number;
            source-address address;
            version (5 | 8);
            version9;
        }
        interface interface-name {
            engine-id number;
            engine-type number;
            source-address address;
        }
    }
```

**Hierarchy Level** [edit **forwarding-options sampling family** (inet | inet6 | mpls)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
**version9** statement introduced in Junos OS 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.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Active Flow Monitoring Using Version 9 on page 220</a></li><li>• <a href="#">Configuring the Output File for Traffic Sampling on page 215</a></li></ul>

## packet-capture

---

<b>Syntax</b>	<pre>packet-capture {     disable;     filefilename filename &lt;files number&gt; &lt;size bytes&gt; &lt;world-readable   no-world-readable&gt;;     maximum-capture-size number; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Configure packet capture on a router.
<b>Options</b>	<b>disable</b> —Disable packet capture on the router.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Packet Capture on page 250</a></li></ul>

## per-flow

---

<b>Syntax</b>	<pre>per-flow {     hash-seed; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options load-balance</a> ], [edit logical-systems <i>logical-system-name</i> <a href="#">forwarding-options load-balance</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <a href="#">forwarding-options load-balance</a> ], [edit routing-instances <i>routing-instance-name</i> <a href="#">forwarding-options load-balance</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3 (M120, M320, and MX Series routers only).
<b>Description</b>	Enable per-flow load balancing based on hash values.
<b>Options</b>	<b>hash-seed</b> —Configure the hash value. The Junos OS automatically chooses a value for the hashing algorithm used. You cannot configure a specific hash value for per-flow load balancing.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Per-Flow Load Balancing Based on Hash Values on page 239</a></li><li>• <a href="#">load-balance (Forwarding Options) on page 345</a></li></ul>



## per-prefix

<b>Syntax</b>	<code>per-prefix {     hash-seed <i>number</i>; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options load-balance</a> ], [edit logical-systems <i>logical-system-name</i> <a href="#">forwarding-options load-balance</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <a href="#">forwarding-options load-balance</a> ], [edit routing-instances <i>routing-instance-name</i> <a href="#">forwarding-options load-balance</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the hash parameter for per-prefix load balancing.
<b>Options</b>	<b>hash-seed</b> —Per-prefix load-balancing hash function.  <b>number</b> —Hash value. <b>Range:</b> 0 through 65,534 <b>Default:</b> 0
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Per-Prefix Load Balancing on page 238</a></li> <li>• <a href="#">load-balance (Forwarding Options) on page 345</a></li> </ul>

## port (cflowd)

<b>Syntax</b>	<code>port <i>port-number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options accounting group-name output cflowd hostname</a> ], [edit <a href="#">forwarding-options monitoring group-name family</a> inet <a href="#">output flow-server hostname</a> ], [edit <a href="#">forwarding-options sampling family</a> (inet   inet6   mpls) <a href="#">output flow-server hostname</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the UDP port number on the cflowd host system.
<b>Options</b>	<b>port-number</b> —Any valid UDP port number on the host system.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Flow Aggregation (cflowd) on page 217</a></li> </ul>

## port (Packet Forwarding)

**Syntax** `port port-number {  
     description text-description;  
     interface interface-name {  
         broadcast;  
         description text-description;  
         no-listen;  
         server address <logical-system logical-system-name> <routing-instance routing-instance-name>;  
     }  
     server address <logical-system logical-system-name> <routing-instance routing-instance-name>;  
}`

**Hierarchy Level** [edit [forwarding-options](#) helpers]

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

**Description** Configure the UDP or TCP port number for listening.

In addition to TFTP and DNS, you can use the helpers statement to enable LAN broadcast forwarding. To do this, you must manually define which UDP port is forwarded, as shown here:

```
user@ host# show forwarding-options
helpers {
  port 3000 {
    interface {
      fe-0/0/1.0 {
        server 111.0.0.2;
      }
    }
  }
  port 3001 {
    interface {
      fe-0/0/0.0 {
        server 100.0.0.2;
      }
    }
  }
}
```

The remaining statements are explained separately.

**Options** *port-number*—UDP or TCP port number for listening.  
**Range:** 1 through 65535

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

- Related Documentation**
- [Configuring DNS and TFTP Packet Forwarding on page 241](#)
  - [Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239](#)

## port-mirroring

```
Syntax  port-mirroring {
        input {
            maximum-packet-length bytes;
            rate number;
            run-length number;
        }
        family (ccc | inet | inet6 | vpls) {
            output {
                interface interface-name {
                    next-hop address;
                }
                no-filter-check;
            }
        }
        instance {
            instance-name {
                input {
                    maximum-packet-length bytes;
                    rate number;
                    run-length number;
                }
                family (ccc | inet | inet6 | vpls) {
                    output {
                        interface interface-name {
                            next-hop address;
                        }
                        no-filter-check;
                    }
                }
            }
        }
        mirror-once;
        traceoptions {
            file filename <files number> <size bytes> <world-readable | no-world-readable>;
        }
    }
```

Hierarchy Level [edit [forwarding-options](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
**family vpls** statement introduced in Junos OS Release 9.3 (MX Series routers only); support extended to M7i, M10, M120, and M320 routers in Junos OS Release 9.5.  
**instance *port-mirroring-instance-name*** statement introduced in Junos OS Release 9.3 (MX Series routers only); support extended to M120 and M320 routers in Junos OS Release 9.5.  
**mirror-once** statement introduced in Junos OS Release 9.3 (MX Series routers only); support extended to M120 routers in Junos OS Release 9.5.  
**family ccc** statement introduced in Junos OS Release 9.6 (M120 and M320 routers only)  
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.

<b>Description</b>	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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Port Mirroring on page 246</a></li><li>• JUNOS® MX Series 3D Universal Edge Routers Solutions, Release 13.1</li></ul>

---

## rate (Forwarding Options)

---

<b>Syntax</b>	<code>rate <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options port-mirroring input</a> ], [edit <a href="#">forwarding-options port-mirroring instance <i>instance-name</i> input</a> ], [edit <a href="#">forwarding-options sampling input</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.
<b>Description</b>	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.
<b>Options</b>	<i>number</i> —Denominator of the ratio. <b>Range:</b> 1 through 65,535
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Port Mirroring on page 246</a></li><li>• <a href="#">Configuring Traffic Sampling on page 213</a></li></ul>

## relay-agent-option

---

<b>Syntax</b>	relay-agent-option;
<b>Hierarchy Level</b>	[edit logical-systems routing-instances <i>instance-name</i> forwarding-options helpers bootp], [edit forwarding-options helpers bootp], [edit routing-instances <i>instance-name</i> forwarding-options helpers bootp],
<b>Release Information</b>	Statement introduced in Junos OS Release 8.0.
<b>Description</b>	Enables the DHCP relay agent information option which allows DHCP to forward information between different VRF routing instances. The functionality is described in RFC 3046, <i>DHCP Relay Agent Information Option</i> . For the Junos OS implementation, the DHCP option number is 82 and the suboption ID is 1. The suboption length is the length required to contain an interface name in addition to the terminating null character. The overall option length is the suboption length plus 2 bytes (for the option header). The DHCP relay agent information option is only present on packets sent between the relay and the DHCP server.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239</a></li></ul>

## route-accounting

---

<b>Syntax</b>	route-accounting;
<b>Hierarchy Level</b>	[edit forwarding-optionsfamily inet6]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3.
<b>Description</b>	Configure the routing platform to track IPv6 traffic passing through the router.
<b>Default</b>	Disabled
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IPv6 Accounting on page 234</a></li></ul>

## run-length

---

<b>Syntax</b>	<code>run-length <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit forwarding-options port-mirroring input], [edit forwarding-options port-mirroring instance <i>port-mirroring-instance-name</i> input], [edit forwarding-options port-mirroring family (inet inet6) input], [edit forwarding-options sampling input], [edit forwarding-options sampling instance <i>instance-name</i> input]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.
<b>Description</b>	Set the number of samples following the initial trigger event. This allows you to sample packets following those already being sampled.
<b>Options</b>	<i>number</i> —Number of samples. <b>Range:</b> 0 through 20 <b>Default:</b> 0
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Applying Filters to Forwarding Tables on page 231</a></li> <li>• <a href="#">Configuring Port Mirroring</a></li> <li>• <a href="#">Configuring Traffic Sampling</a></li> </ul>

## sampling (Forwarding Options)

```

Syntax  sampling {
        disable;
        sample-once;
        family (inet | inet6 | mpls) {
            disable;
            output {
                aggregate-export-interval seconds;
                extension-service service-name;
                file {
                    disable;
                    filename filename;
                    files number;
                    size bytes;
                    (stamp | no-stamp);
                    (world-readable | no-world-readable);
                }
                flow-active-timeout seconds;
                flow-inactive-timeout seconds;
                flow-server 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;
                    version9 {
                        template template-name;
                    }
                }
                interface interface-name {
                    engine-id number;
                    engine-type number;
                    source-address address;
                }
            }
        }
    }
    input {
        max-packets-per-second number;
        maximum-packet-length bytes;
        rate number;
        run-length number;
    }
    instance instance-name {
        disable;
    }

```



```

family (inet | inet6 | mpls) {
  disable;
  output {
    aggregate-export-interval seconds;
    extension-service service-name;
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    flow-server 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;
      version-ipfix {
        template template-name;
      }
      version9 {
        template template-name;
      }
    }
    inline-jflow {
      source-address address;
      flow-export-rate rate;
    }
    interface interface-name {
      engine-id number;
      engine-type number;
      source-address address;
    }
  }
}
input {
  max-packets-per-second number;
  maximum-packet-length bytes;
  rate number;
  run-length number;
}
}
traceoptions {
  no-remote-trace;
  file filename <files number> <size bytes> <match expression> <world-readable |
  no-world-readable>;
}
}

```

Hierarchy Level [edit forwarding-options]

<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure traffic sampling.  The statements are explained separately.
<b>Usage Guidelines</b>	See Configuring Traffic Sampling.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Applying Filters to Forwarding Tables on page 231</a></li><li>• <a href="#">Configuring Active Flow Monitoring Using Version 9 on page 220</a></li><li>• <a href="#">Configuring Flow Aggregation (cflowd) on page 217</a></li><li>• <a href="#">Configuring Port Mirroring on page 246</a></li><li>• <a href="#">Tracing Traffic-Sampling Operations on page 217</a></li></ul>

---

## server (DHCP and BOOTP Relay Agent)

---

<b>Syntax</b>	<pre>server <i>address</i> {     logical-system <i>logical-system-name</i> &lt;routing-instance [ &lt;default&gt;         <i>routing-instance-names</i> ]&gt;;     routing-instance [ &lt;default&gt; <i>routing-instance-names</i> ]; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpers bootp</a> ], [edit <a href="#">forwarding-options helpers bootp interface</a> ( <i>interface-name</i>   <i>interface-group</i> )]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Configure the router or switch to act as a DHCP and BOOTP relay agent.
<b>Options</b>	<ul style="list-style-type: none"><li>• <b>address</b>—One or more addresses of the server.</li><li>• <b>logical-system <i>logical-system-name</i></b>—(Optional) Logical system of the server.</li><li>• <b>routing-instance <i>routing-instance-names</i></b>—(Optional) Routing instance name that belong to the DHCP or BOOTP relay agent.</li></ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents on page 239</a></li></ul>

## server (DNS and TFTP Service)

<b>Syntax</b>	<code>server address &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpers domain</a> ], [edit <a href="#">forwarding-options helpers domain interface interface-name</a> ], [edit <a href="#">forwarding-options helpers tftp</a> ], [edit <a href="#">forwarding-options helpers tftp interface interface-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the DNS or TFTP server for forwarding DNS or TFTP requests. Only one server can be specified for each interface.
<b>Options</b>	<b>address</b> —Address of the server.  <b>logical-system logical-system-name</b> —(Optional) Logical system of the server.  <b>routing-instance [ routing-instance-names ]</b> —(Optional) Set the routing instance name or names that belong to the DNS server or TFTP server.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring DNS and TFTP Packet Forwarding on page 241</a></li> </ul>

## size (Packet Capture)

<b>Syntax</b>	<code>size number;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options packet-capture file</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Configure the maximum size of the file for packet capturing.
<b>Options</b>	<b>number</b> —Maximum size of file. <b>Range:</b> 1024 through 104,857,600 bytes <b>Default:</b> 512,000 bytes
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Packet Capture on page 250</a></li> </ul>

## size (Sampling and Traceoptions)

---

<b>Syntax</b>	<code>size bytes;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpertraceoptions file</a> ], [edit <a href="#">forwarding-options port-mirroring traceoptions file</a> ], [edit <a href="#">forwarding-options sampling family family-name output file</a> ], [edit <a href="#">forwarding-options sampling traceoptions file</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.
<b>Description</b>	<p>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.</p> <p>When a traffic sampling file named <b>sampling-file</b> reaches the maximum size, it is renamed <b>sampling-file.0</b>. When the <b>sampling-file</b> file again reaches its maximum size, <b>sampling-file.0</b> is renamed <b>sampling-file.1</b> and <b>sampling-file</b> is renamed <b>sampling-file.0</b>. This renaming scheme continues until the maximum number of traffic sampling files is reached. Then the oldest traffic sampling file is overwritten.</p>
<b>Options</b>	<p><b>bytes</b>—Maximum size of each traffic sampling file or trace log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).</p> <p><b>Syntax:</b> <code>xk</code> to specify KB, <code>xm</code> to specify MB, or <code>xg</code> to specify GB</p> <p><b>Range:</b> 10 KB through the maximum file size supported on your router</p> <p><b>Default:</b> 1 MB for sampling data; 128 KB for log information</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Output File for Traffic Sampling on page 215</a></li><li>• <a href="#">Tracing Traffic-Sampling Operations on page 217</a></li></ul>

## stamp

<b>Syntax</b>	(stamp   no-stamp);
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options sampling family family-name output file</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Include a timestamp with each line in the output file.
<b>Default</b>	no-stamp
<b>Options</b>	<p><b>no-stamp</b>—Do not include timestamps.</p> <p><b>stamp</b>—Include a timestamp with each line of packet sampling information.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Output File for Traffic Sampling on page 215</a></li> </ul>

## tftp

<b>Syntax</b>	<pre>tftp {   <b>description</b> text-description;   <b>interface</b> interface-name {     broadcast;     <b>description</b> text-description;     no-listen;     <b>server</b> address &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;;   }   <b>server</b> address &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpers</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Enable TFTP request packet forwarding.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring DNS and TFTP Packet Forwarding on page 241</a></li> </ul>

## traceoptions (DNS and TFTP Packet Forwarding)

<b>Syntax</b>	<pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;size <i>bytes</i>&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i>;     level <i>level</i>;     &lt;no-remote-trace&gt;; } </pre>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helpers</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement standardized and <b>match</b> option introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Configure tracing operations for BOOTP, DNS and TFTP packet forwarding.
<b>Default</b>	If you do not include this statement, no tracing operations are performed.
<b>Options</b>	<p><b>file <i>filename</i></b>—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 <b>fud</b> in the directory <b>/var/log</b>. If you include the <b>file</b> statement, you must specify a filename.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, 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 <b>size</b> option and a filename.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. You can include the following flags:</p> <ul style="list-style-type: none"> <li>• <b>address</b>—Trace address management events</li> <li>• <b>all</b>—Trace all events</li> <li>• <b>bootp</b>—Trace BOOTP or DHCP services events</li> <li>• <b>config</b>—Trace configuration events</li> <li>• <b>domain</b>—Trace DNS service events</li> <li>• <b>ifdb</b>—Trace interface database operations</li> <li>• <b>io</b>—Trace I/O operations</li> <li>• <b>main</b>—Trace main loop events</li> <li>• <b>port</b>—Trace arbitrary protocol events</li> </ul>

- **rtsock**—Trace routing socket operations
- **tftp**—Trace TFTP service events
- **trace**—Trace tracing operations
- **ui**—Trace user interface operations
- **util**—Trace miscellaneous utility operations

**match *regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

**no-remote-trace**—(Optional) Disable remote tracing globally or for a specific tracing operation.

**no-world-readable**—(Optional) Restrict file access to the owner.

**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.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tracing BOOTP, DNS, and TFTP Forwarding Operations on page 242</a></li></ul>

## traceoptions (Port Mirroring and Traffic Sampling)

---

<b>Syntax</b>	<code>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>bytes</i>&gt; &lt;world-readable   no-world-readable&gt;; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options port-mirroring</a> ], [edit <a href="#">forwarding-options sampling</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure traffic sampling tracing operations.  The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tracing Traffic-Sampling Operations on page 217</a></li></ul>

## version

---

<b>Syntax</b>	<code>version <i>format</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options accounting group-name output cflowd hostname</a> ], [edit <a href="#">forwarding-options sampling family</a> (inet   inet6   mpls) <a href="#">output flow-server hostname</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the version format of the aggregated flows exported to a cflowd server.
<b>Options</b>	<i>format</i> —Export format of the flows. <b>Values:</b> 5 or 8 <b>Default:</b> 5
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flow Aggregation (cflowd) on page 217</a></li></ul>



## version9

---

<b>Syntax</b>	<code>version9 {     template <i>template-name</i>; }</code>
<b>Hierarchy Level</b>	[edit <b>forwarding-options sampling family</b> <i>family-name</i> <b>output flow-server</b> <i>hostname</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3.
<b>Description</b>	Enable active flow monitoring using the version 9 template format to collect traffic flows.
<b>Options</b>	<b>template <i>template-name</i></b> —Name of a version 9 record flow format template configured at the [edit <b>services monitoring</b> ] hierarchy level.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Active Flow Monitoring Using Version 9 on page 220</a></li><li>• Flow Monitoring</li><li>• Junos Services Interfaces Configuration Release 12.3</li></ul>

## world-readable (Forwarding Options)

---

<b>Syntax</b>	(world-readable   no-world-readable);
<b>Hierarchy Level</b>	[edit <a href="#">forwarding-options helperstraceoptions file</a> ], [edit <a href="#">forwarding-options packet-capture file</a> ], [edit <a href="#">forwarding-options port-mirroring traceoptions file</a> ], [edit <a href="#">forwarding-options sampling family <i>family-name</i> output file</a> ], [edit <a href="#">forwarding-options samplingtraceoptions file</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.
<b>Description</b>	Enable unrestricted file access.
<b>Default</b>	no-world-readable
<b>Options</b>	<b>no-world-readable</b> —Restrict file access to the owner.  <b>world-readable</b> —Enable unrestricted file access.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Output File for Traffic Sampling on page 215</a></li><li>• <a href="#">Tracing Traffic-Sampling Operations on page 217</a></li></ul>

## PART 4

# Index

- [Index on page 383](#)



# Index

## Symbols

!	
in policy expressions	
logical operator.....	72
#, comments in configuration statements.....	xx
&&, logical operator.....	71
( ), in syntax descriptions.....	xx
< >, in syntax descriptions.....	xx
[ ], in configuration statements.....	xx
{ }, in configuration statements.....	xx
(pipe), in syntax descriptions.....	xx
(pipes), logical operator.....	71

## A

accept	
firewall filters	
action.....	212
policy, routing	
control action.....	60
access and access-internal routes	
importing and exporting in routing policies.....	98
accounting statement.....	298
usage guidelines.....	234
actions	
default, routing policy.....	23
final, routing policy.....	23
policy, routing.....	23
characteristics, manipulating.....	61
flow control.....	58, 60
tracing.....	58
route list match types.....	22
routing policy, summary of.....	9
tracing.....	69
address-mask (route list match type).....	158
evaluation in a route list.....	161
example.....	167, 168, 169
aggregated Ethernet	
load balancing.....	197
aggregation statement.....	299
usage guidelines.....	234

aigp statement	
BGP.....	264
apply-path statement.....	265
usage guidelines.....	152
area (routing policy match condition).....	52
AS path	
prepending.....	178
as-path (routing policy match condition).....	53
as-path statement.....	266
policy, routing	
usage guidelines.....	134
as-path-group statement.....	267
usage guidelines.....	134
as-path-prepend (routing policy action).....	61, 177
ASs	
paths	
modifying with routing policy.....	61, 177
regular expressions See policy, routing, AS	
path regular expressions	
autonomous-system-type statement.....	300
usage guidelines.....	217

## B

BGP	
communities	
names.....	140
policy, routing.....	31, 268
damping parameters.....	181, 271
dynamic routing policies	
applying.....	79
extended communities.....	144
policy, routing.....	273, 282
applying.....	261
proxy route target filtering.....	295
BGP (Border Gateway Protocol)	
policy to make routes less preferable.....	178
BOOTP relay agent.....	239
bootp statement.....	301
usage guidelines.....	239
braces, in configuration statements.....	xx
brackets	
angle, in syntax descriptions.....	xx
square, in configuration statements.....	xx
BSR	
policy, import.....	287

## C

class (routing policy action).....	61
client-response-ttl statement.....	303

color	
policy, routing	
action.....	61
match condition.....	53
comments, in configuration statements.....	xx
communities	
extend range of BGP communities.....	144
names.....	140
policy, routing.....	31, 268
action.....	61
match condition.....	53
community statement.....	268
policy, routing	
usage guidelines.....	140
condition statement.....	270
conventions	
text and syntax.....	xix
curly braces, in configuration statements.....	xx
customer support.....	xxi
contacting JTAC.....	xxi
<b>D</b>	
damping	
policy, routing, action.....	62
damping statement.....	271
BGP	
usage guidelines.....	184
policy, routing	
usage guidelines.....	181
default routing policy See PTX Series	
defaults	
routing policy actions.....	23
defaults statement	
static statement	
usage guidelines.....	88
description statement	
helper service or interface.....	304
service	
usage guidelines.....	239, 241
destination class usage.....	62, 63
destination-class (routing policy action).....	62, 66
DHCP	
relay agent.....	239
relay agents, extended	
access and access-internal routes.....	98
snoping.....	245
dhcp-relay statement	
DHCP snooping.....	305
disable statement	
packet capture.....	305
sampling.....	305
traffic sampling	
usage guidelines.....	215
discard interface.....	81
discard statement, in static statement	
usage guidelines.....	88
DNS	
packet forwarding.....	241
requests, disabling recognition.....	241
documentation	
comments on.....	xxi
domain statement.....	306
DVMRP	
policy, routing.....	275, 283
applying.....	261
dynamic database	
active nonstop routing.....	79
routing policies.....	77
dynamic routing policies	
active nonstop routing.....	79
BGP.....	79
configuring.....	77
dynamic-db statement.....	272
overview.....	76
dynamic-db statement.....	272
usage guidelines.....	77
<b>E</b>	
ECMP	
flow-based forwarding.....	126, 127
enhanced-hash-key statement	
forwarding-options.....	307
equal-cost multi-path See ECMP	
ether-pseudowire statement.....	317
usage guidelines.....	193
evaluation	
policy, routing .....	26
exact (route list match type).....	158
exact route list match type.....	22, 154
export routing policies	
applying.....	24, 70
export statement.....	275
BGP.....	273
DVMRP.....	275
forwarding table.....	281
IS-IS.....	274
MSDP.....	276

OSPF.....	277	in traffic sampling.....	212
PIM.....	278	purpose.....	11
policy, routing		flood statement.....	327
usage guidelines.....	24, 70, 75	usage guidelines.....	231
RIP.....	279	flooding	
RIPng.....	280	IS-IS and OSPF.....	19
export-format statement.....	306	flow aggregation.....	217
usage guidelines.....	235	flow control actions.....	58, 60
		flow control, actions in routing policies.....	9
		flow-active-timeout statement.....	327
		accounting	
		usage guidelines.....	234
		sampling	
		usage guidelines.....	215
		flow-based forwarding	
		ECMP.....	126, 127
		flow-export-destination statement.....	328
		usage guidelines.....	235
		flow-inactive-timeout statement.....	328
		accounting	
		usage guidelines.....	234
		sampling	
		usage guidelines.....	215
		flow-server statement	
		usage guidelines.....	217
		font conventions.....	xix
		forwarding classes	
		policy to group source and destination	
		prefixes.....	106
		forwarding table	
		filters.....	229, 231
		policy, routing.....	281
		applying.....	25, 75
		static routes.....	88
		forwarding-options statement.....	330
		usage guidelines.....	225
		from statement.....	291
		policy, routing	
		usage guidelines.....	51
		FTP traffic, sampling.....	223
		<b>G</b>	
		group statement	
		DHCP snooping.....	330
		<b>H</b>	
		hash-key statement.....	331
		usage guidelines.....	225
<b>F</b>			
family inet statement (load balancing).....	315		
usage guidelines.....	225		
family mpls statement.....	317		
usage guidelines.....	189, 193		
family multiservice statement.....	320		
MX Series routers			
usage guidelines.....	196		
usage guidelines.....	194		
VPLS load balancing			
usage guidelines.....	195		
family statement			
forwarding table filters.....	310		
port mirroring.....	312		
sampling.....	313		
file statement			
helpers trace options.....	322		
packet capture.....	322		
sampling.....	323		
traceoptions.....	323		
traffic sampling output			
usage guidelines.....	215, 217, 242		
files			
logging information output file.....	217, 242		
traffic sampling output files.....	215		
var/log/sampled file.....	217, 242		
var/tmp/sampled.pkts file.....	215		
files statement			
packet capture.....	325		
sampling.....	325		
usage guidelines.....	215		
filter statement			
forwarding table.....	326		
VPLS.....	326		
filtering received labels.....	283		
firewall filters			
applying.....	229		
architecture .....	6, 11		
comparison with routing policies .....	11, 14		
flow, packets .....	4		

hash-seed statement	
load balancing.....	362, 363
helpers statement.....	333
<b>I</b>	
if-route-exists statement.....	270
import routing policies	
applying.....	24, 70, 71
overview.....	17
import statement	
BGP.....	282
bootstrap.....	287
DVMRP.....	283
LDP.....	283
MSDP.....	284
OSPF.....	285
PIM.....	286
policy, routing	
usage guidelines.....	24, 70
RIP.....	288
RIPng.....	289
route resolution.....	290
indexed-load-balance.....	335
input statement	
firewall filters	
usage guidelines.....	229
forwarding table.....	336
port mirroring.....	336
sampling.....	337
traffic sampling	
usage guidelines.....	213
usage guidelines.....	231
install-nexthop lsp (routing policy action).....	62
install-to-fib	
usage guidelines.....	123
instance (routing policy match condition).....	54
instance statement	
port mirroring.....	338
instance-name.inet.0 routing table.....	256
interface (routing policy match condition).....	54
interface statement	
accounting or sampling.....	339
BOOTP.....	340
DNS or TFTP packet forwarding or relay	
agent.....	342
monitoring.....	343
next-hop group.....	344
port mirroring.....	344
snooping.....	341
usage guidelines.....	239
invert-match statement	
usage guidelines.....	146
IP addresses	
sampling traffic from single IP addresses.....	222
IPv6	
static routes.....	93
IPv6 accounting, configuring.....	234
IS-IS	
policy, routing.....	274
applying.....	261
<b>J</b>	
joins, PIM	
rejecting.....	166
<b>L</b>	
label filtering.....	283
layer-3 statement	
load balancing	
usage guidelines.....	196
layer-3-only statement	
usage guidelines.....	195
layer-4 statement	
load balancing	
usage guidelines.....	196
LDP	
policy filters.....	283
policy, routing	
applying.....	261
received label filtering.....	283
level (routing policy match condition).....	54
load balancing	
configuring.....	197
Ethernet pseudowires.....	193
MPLS LSPs.....	189
per-flow.....	41, 239
per-packet.....	37
IPv4.....	186
per-prefix.....	41, 123, 238
VPLS	
M120 and M320 routers.....	195
MX Series routers.....	196
load-balance group.....	237
load-balance statement.....	345
usage guidelines.....	238, 239
load-balance-group statement.....	347
usage guidelines.....	237



local-dump statement.....	347
usage guidelines.....	217, 219
local-preference	
policy, routing	
action.....	63
match condition.....	54
log output	
traffic sampling.....	217, 242
longer (route list match type).....	158
longer route list match type.....	22, 155
LSPs	
load balancing.....	189

## M

manuals	
comments on.....	xxi
match conditions	
policy, routing.....	51, 258
routing policy, summary of.....	260
match types.....	22
max-packets-per-second statement.....	348
maximum-capture-size statement.....	348
maximum-hop-count statement.....	349
usage guidelines.....	239
maximum-packet-length statement.....	350
metric	
policy, routing	
action.....	63
match condition.....	54
minimum-wait-time statement.....	351
mirror-once statement.....	351
monitoring statement.....	352
usage guidelines.....	235
MPLS	
load balancing.....	189
policy, routing	
applying.....	261
MSDP	
policy, routing.....	276, 284
multicast-scoping	
policy, routing	
match condition.....	55

## N

names	
policy, routing.....	50
neighbor (routing policy match condition).....	55
next policy (routing policy control action).....	60
next term (routing policy control action).....	60

next-hop	
policy, routing	
action.....	65
match condition.....	55
next-hop groups.....	237
next-hop statement.....	353
next-hop groups	
usage guidelines.....	237
next-hop-group statement.....	354
usage guidelines.....	237
next-table statement	
usage guidelines.....	88
no-filter-check statement.....	355
no-install-to-fib	
usage guidelines.....	123
no-listen statement	
usage guidelines.....	239
no-local-dump statement.....	347
usage guidelines.....	217
no-stamp statement.....	375
usage guidelines.....	215
no-world-readable statement.....	380
usage guidelines.....	217

## O

origin	
policy, routing	
action.....	65
match condition.....	56
orlonger (route list match type).....	159
orlonger route list match type.....	22, 155
OSPF	
policy, routing .....	277, 285
applying.....	261
output files	
logging information output file.....	217, 242
traffic sampling output files.....	215
output statement	
accounting.....	356
firewall filters	
usage guidelines.....	229
forwarding table.....	357
monitoring.....	358
port mirroring.....	359
sampling.....	360
usage guidelines.....	231

## P

P2MP.....	197
-----------	-----

packet capture.....	250	IS-IS.....	274
packet-capture statement.....	361	match conditions.....	51, 58, 258
parentheses, in syntax descriptions.....	xx	MSDP.....	276, 284
payload statement		multiple policies	
usage guidelines.....	195	applying.....	70
per-flow load balancing.....	239	evaluation.....	27
per-flow statement.....	362	OSPF.....	277, 285
per-prefix load balancing.....	123, 238	overview.....	15
per-prefix statement.....	363	PIM.....	286
permanent routes, adding.....	88	policy expressions .....	70, 75
PIM		preferences, modifying.....	66
multicast traffic joins, rejecting.....	166	prefix list .....	151, 293
policy, routing.....	286	prefix list filter.....	154, 294
applying.....	261	purpose.....	11
policy (routing policy match condition).....	56	rejecting PIM multicast traffic joins.....	166
policy filters, LDP.....	283	RIP.....	279, 288
policy framework		RIPng.....	280, 289
architecture .....	6	route lists.....	155, 166
comparison of policies .....	11	route target prefix list.....	295
firewall filters.....	3	source prefixes, group.....	105
overview.....	3	subroutines .....	28, 170, 173
policy, routing.....	3	terms.....	50
policy, import		testing.....	30
BSR.....	287	uses for.....	20
policy, routing		policy-options statement.....	290
access and access-internal routes.....	98	usage guidelines.....	45
actions.....	23, 58, 62, 64, 69	policy-statement statement.....	291
applying.....	70	from statement.....	51
overview.....	24	then statement.....	58
architecture .....	6	to statement.....	51
AS path regular		usage guidelines.....	50
expressions.....	133, 140, 265, 266, 267	port mirroring.....	246
BGP.....	273, 282	mirror-once statement.....	351
BGP damping parameters.....	271	multiple instances.....	338
chains		port statement.....	363, 364
applying.....	70	usage guidelines.....	217
evaluation.....	27	port-mirroring statement.....	366
communities.....	31, 268	usage guidelines.....	246
comparison with firewall filters .....	11	preferences	
configuring.....	46, 47, 83	modifying	
overview .....	20, 25	with routing policies.....	66
default policies and actions.....	18	policy, routing	
DVMRP.....	275, 283	action.....	66
evaluation .....	26	match condition.....	56
export policies.....	24	prefix list .....	151, 293
flow, routing information .....	4	prefix list filter.....	294
forwarding table.....	281	prefix-length-range (route list match type).....	159
framework.....	15	prefix-length-range match type.....	22
import policies.....	17, 24, 75	prefix-list (routing policy match condition).....	56

- prefix-list statement.....293
  - usage guidelines.....152
- prefix-list-filter statement.....294
- protocols
  - applying policies.....24
  - match condition
    - policy, routing.....56
  - routing
    - applying policies.....75
- PTX Series
  - default routing policy.....123
- R**
  - rate statement.....367
    - usage guidelines.....213
  - receive routes.....88
  - receive statement
    - static routes
      - usage guidelines.....88
  - received label filtering.....283
  - reject
    - policy, routing
      - control action.....60
  - reject option to static statement
    - usage guidelines.....88
  - relay agents
    - DHCP and BOOTP.....239
  - relay agents, DHCP
    - extended
      - access and access-internal routes.....98
  - relay-agent-option statement.....368
    - usage guidelines.....239
  - rib (routing policy match condition).....57
  - RIP
    - policy, routing.....279, 288
      - applying.....262
  - RIPng
    - policy, routing.....280, 289
  - route list match types.....22
  - route lists.....155
  - route manipulation actions, routing policies.....9
  - route recording.....217
  - route statement
    - static statement
      - usage guidelines.....88, 93
  - route target prefix list.....295
  - route-accounting statement.....368
    - usage guidelines.....234
  - route-filter (routing policy match condition).....57
  - routers
    - DHCP relay agents.....239
  - routes
    - static See static routes
  - routing
    - static See static routing
  - routing policies
    - configuration tasks.....83, 85, 106, 127, 178
    - default actions.....23
    - dynamic
      - configuring.....77
    - dynamic database.....77
    - ECMP flow-based forwarding.....126, 127
    - final actions.....23
    - forwarding class with source and
      - destination.....106
    - grouping source and destination prefixes.....106
    - making BGP routes less preferable.....178
    - policy name.....83
    - preparation.....10
    - prepending AS paths.....178
    - terms.....23
    - terms, creating.....85
  - routing policy See policy, routing
  - routing solutions
    - making BGP routes less preferable.....178
  - routing tables
    - instance-name.init.0.....256
  - routing-instance statement
    - usage guidelines.....239
  - rtf-prefix-list statement.....295
  - run-length statement.....369
    - usage guidelines.....213
- S**
  - sample (firewall filter action).....212
  - sampled file.....217, 242
  - sampled.pkts file.....215
  - sampling statement
    - usage guidelines.....211, 212
  - server statement
    - DHCP and BOOTP service .....372
    - DNS and TFTP service .....373
    - usage guidelines.....239
  - show chassis hardware command
    - usage guidelines.....39
  - show policy damping command.....186
    - usage guidelines.....181

show route detail command	
usage guidelines.....	181
show route receive-protocol command	
usage guidelines.....	58
size statement	
packet capture.....	373
sampling.....	374
snooping, DHCP.....	245
SONET interfaces	
sampling.....	221
source class usage.....	66, 105
source-address-filter (routing policy match condition).....	57
source-class (routing policy action).....	66
stamp option.....	216
stamp statement.....	375
usage guidelines.....	215
static routes.....	88, 93
configuring basic routes (configuration editor).....	88
IPv6.....	93
static routing	
overview.....	88
<i>See also</i> static routes	
static statement	
usage guidelines.....	88, 93
subroutines .....	28, 170
support, technical <i>See</i> technical support	
syntax conventions.....	xix
<b>T</b>	
table statement.....	270
tag	
policy, routing	
action.....	67
technical support	
contacting JTAC.....	xxi
term statement	
policy	
usage guidelines.....	50
terms	
in a routing policy.....	23
in a routing policy, creating.....	85
policy, routing.....	50
test policy command .....	30
TFTP	
packet forwarding.....	241
requests, disabling recognition.....	241
tftp statement.....	375
then statement.....	291
policy, routing	
usage guidelines.....	58
through (route list match type).....	159
through route list match type.....	22
timestamp option.....	216
to statement.....	291
usage guidelines.....	51
trace (policy tracing action).....	58, 69
traceoptions statement	
DNS and TFTP packet forwarding.....	376
port mirroring and traffic sampling.....	378
usage guidelines.....	228
tracing actions.....	58, 69
traffic	
accounting.....	234
forwarding	
configuration statements.....	225
overview.....	225
monitoring.....	235
sampling	
configuration statements.....	211
disabling.....	215, 305
DNS and TFTP packet forwarding.....	241
example configurations.....	221
flow aggregation.....	217
FTP traffic.....	223
logging information output	
file.....	217, 231, 242
output files.....	215
overview.....	211
run-length parameter.....	213
sampling rate parameter.....	213
SONET interfaces.....	221
traffic from single IP addresses.....	222
traffic sampling	
configuring.....	212
<b>U</b>	
upto (route list match type).....	159
upto route list match type.....	22
<b>V</b>	
var/log/sampled file.....	217, 231, 242
var/tmp/sampled.pkts file.....	215
verification	
static route.....	91, 96
version statement.....	378
usage guidelines.....	217

version9 statement.....	379
usage guidelines.....	220
VPLS.....	197
VPLS load balancing	
M120 and M320 routers.....	195
MX Series routers.....	196
 <b>W</b>	
world-readable statement.....	380
usage guidelines.....	217

