




Junos[®] OS for EX Series Ethernet Switches, Release 12.3: Layer 3 Protocols



Published: 2013-01-04
Revision 1

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
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 for EX Series Ethernet Switches, Release 12.3: Layer 3 Protocols

Copyright © 2013, Juniper Networks, Inc.

All rights reserved.

Revision History

January 2013—Revision 1

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

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 This Topic Collection	xv
	How to Use This Guide	xv
	List of EX Series Guides for Junos OS Release 12.3	xv
	Downloading Software	xviii
	Documentation Symbols Key	xviii
	Documentation Feedback	xx
	Requesting Technical Support	xx
	Self-Help Online Tools and Resources	xx
	Opening a Case with JTAC	xxi
Part 1	Layer 3 Protocols	
Chapter 1	Layer 3 Protocols—Overview	3
	Layer 3 Protocols Supported on EX Series Switches	3
	Layer 3 Protocols Not Supported on EX Series Switches	4
	Understanding Distributed Periodic Packet Management on EX Series Switches	6
	Understanding IPsec Authentication for OSPF Packets on EX Series Switches	7
	Authentication Algorithms	7
	Encryption Algorithms	8
	IPsec Protocols	8
	Security Associations	9
	IPsec Modes	9
Chapter 2	Configuring Layer 3 Protocols	11
	Configuring BGP Sessions (J-Web Procedure)	11
	Configuring an OSPF Network (J-Web Procedure)	15
	Configuring a RIP Network (J-Web Procedure)	19
	Configuring Static Routing (CLI Procedure)	24
	Configuring Static Routing (J-Web Procedure)	24
	Configuring Routing Policies (J-Web Procedure)	26
	Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)	31
	Disabling or Enabling Distributed Periodic Packet Management Globally	31
	Disabling or Enabling Distributed Periodic Packet Management for LACP Packets	32
	Using IPsec to Secure OSPFv3 Networks (CLI Procedure)	32
	Configuring Security Associations	32
	Securing OPSFv3 Networks	33

Chapter 3	Verifying Layer 3 Protocols Configuration	35
	Monitoring BGP Routing Information	35
	Monitoring OSPF Routing Information	37
	Monitoring RIP Routing Information	40
	Monitoring Routing Information	41
Chapter 4	Configuration Statements for Layer 3 Protocols	45
	[edit protocols bgp] Configuration Statement Hierarchy on EX Series	
	Switches	45
	Supported Statements in the [edit protocols bgp] Hierarchy Level	46
	Unsupported Statements in the [edit protocols bgp] Hierarchy Level	54
	[edit protocols isis] Configuration Statement Hierarchy on EX Series	
	Switches	55
	Supported Statements in the [edit protocols isis] Hierarchy Level	55
	Unsupported Statements in the [edit protocols isis] Hierarchy Level	57
	[edit protocols ospf] Configuration Statement Hierarchy on EX Series	
	Switches	58
	Supported Statements in the [edit protocols ospf] Hierarchy Level	58
	Unsupported Statements in the [edit protocols ospf] Hierarchy Level	61
	[edit protocols ospf3] Configuration Statement Hierarchy on EX Series	
	Switches	61
	Supported Statements in the [edit protocols ospf3] Hierarchy Level	62
	Unsupported Statements in the [edit protocols ospf3] Hierarchy Level	64
	[edit protocols rip] Configuration Statement Hierarchy on EX Series	
	Switches	65
	Supported Statements in the [edit protocols rip] Hierarchy Level	65
	Unsupported Statements in the [edit protocols rip] Hierarchy Level	67
	[edit protocols ripng] Configuration Statement Hierarchy on EX Series	
	Switches	67
	Supported Statements in the [edit protocols ripng] Hierarchy Level	67
	Unsupported Statements in the [edit protocols ripng] Hierarchy Level	68
	accept-remote-nexthop	69
	active	70
	advertise-external	71
	advertise-inactive	72
	advertise-peer-as	73
	aggregate (Routing)	74
	aggregate-label	75
	allow	76
	any-sender	77
	area	78
	area-range	80
	as-override	82
	as-path (Routing Options)	83
	asm-override-ssm	85
	authentication-algorithm	86
	authentication-key (Protocols BGP)	87
	authentication-key (Protocols IS-IS)	88
	authentication-key (Protocols RIP)	89

authentication-key-chain (Protocols BGP)	90
authentication-key-chains	91
authentication-type (Protocols IS-IS)	92
authentication-type (Protocols RIP)	93
autonomous-system	94
backup-pe-group	96
backups	97
bandwidth (Multicast Flow Map)	98
bandwidth-based-metrics	99
bfd-liveness-detection (Protocols BGP)	101
bfd-liveness-detection (Protocols IS-IS)	105
bfd-liveness-detection (Protocols OSPF)	107
bfd-liveness-detection (Protocols RIP)	111
bfd-liveness-detection (Routing Options Static Route)	114
bgp	118
bgp-orf-cisco-mode	119
bmp	120
brief	121
centralized	122
check-zero	123
checksum (Protocols IS-IS)	124
cluster	125
community (Routing Options)	127
confederation	129
csnp-interval	130
damping (Protocols BGP)	131
dead-interval	133
default-lsa	134
default-metric	135
description (Protocols BGP)	136
disable (Protocols BGP)	137
disable (Protocols IS-IS)	138
disable (OSPF)	139
disable (Routing Options)	140
discard	141
domain-id	142
domain-vpn-tag	143
explicit-null (Protocols BGP)	144
export (Protocols BGP)	145
export (Protocols IS-IS)	146
export (Protocols OSPF)	147
export (Protocols RIP)	148
export (Protocols RIPv6)	149
export (Routing Options)	150
export-rib	151
external-preference (Protocols IS-IS)	152
external-preference (Protocols OSPF)	153
family (Protocols BGP)	154
fate-sharing	158

flow	159
flow-map	160
forwarding-cache (Flow Maps)	161
forwarding-cache (Multicast)	162
forwarding-table	163
generate	164
graceful-restart (Protocols BGP)	166
graceful-restart (Protocols IS-IS)	168
graceful-restart (Protocols OSPF)	169
graceful-restart (Protocols RIP)	171
graceful-restart (Protocols RIPng)	172
graceful-restart (Enabling Globally)	173
group (Protocols BGP)	174
group (Protocols RIP)	177
group (Protocols RIPng)	179
hello-authentication-key	180
hello-authentication-type	181
hello-interval (Protocols IS-IS)	182
hello-interval (Protocols OSPF)	183
hello-padding	184
holddown (Protocols RIP)	185
holddown (Protocols RIPng)	186
hold-time	186
hold-time (Protocols BGP)	187
hold-time (Protocols IS-IS)	189
idle-after-switch-over	190
ignore-attached-bit	191
ignore-lsp-metrics	192
import (Protocols BGP)	193
import (Protocols OSPF)	194
import (Protocols RIP)	195
import (Protocols RIPng)	196
import (Routing Options)	197
import-policy	198
import-rib	199
include-mp-next-hop	200
indirect-next-hop	201
install (Routing Options)	202
instance-export	203
instance-import	203
inter-area-prefix-export	204
inter-area-prefix-import	205
interface (Protocols IS-IS)	206
interface (Protocols OSPF)	208
interface (Routing Options)	210
interface (Multicast Static Routes)	211
interface-routes	212
interface-type (Protocols OSPF)	214
ipv4-multicast	215

ipv4-multicast-metric	216
ipv6-multicast	216
ipv6-multicast-metric	217
ipv6-unicast	218
ipv6-unicast-metric	219
isis	220
keep	221
labeled-unicast (Protocols BGP)	222
level (Global IS-IS)	224
link-protection (Protocols IS-IS)	225
local-address (Protocols BGP)	226
local-address (Routing Options)	228
local-as	229
local-interface (IPv6)	231
local-preference	232
log-updown (Protocols BGP)	233
loose-authentication-check	234
lsp-interval	235
lsp-lifetime	236
lsp-metric-into-summary	237
martians	238
max-areas	239
maximum-bandwidth (Routing Options)	240
maximum-paths	241
maximum-prefixes	243
med-igp-update-interval	244
mesh-group (Protocols IS-IS)	245
message-size	246
metric (Protocols IS-IS)	247
metric (Protocols OSPF Interface)	249
metric (Aggregate, Generated, or Static Route)	250
metric-in (Protocols RIP)	251
metric-in (Protocols RIPng)	252
metric-out (Protocols BGP)	253
metric-out (Protocols RIP)	255
metric-out (Protocols RIPng)	256
metric-type	257
msdp	258
mtu-discovery	260
multicast (Routing Options)	262
multihop	263
multipath (Protocols BGP)	265
neighbor (Protocols BGP)	266
neighbor (Protocols RIP)	269
neighbor (Protocols RIPng)	270
no-adjacency-holddown	271
no-aggregator-id	272
no-authentication-check	273
no-client-reflect	274

no-csnp-authentication	275
no-eligible-backup (Protocols IS-IS)	275
no-hello-authentication	276
no-ipv4-multicast	276
no-ipv4-routing	277
no-ipv6-multicast	278
no-ipv6-routing	279
no-ipv6-unicast	280
no-nssa-abr	281
no-psnp-authentication	282
no-qos-adjust	283
no-rfc-1583	284
no-unicast-topology	285
no-validate	286
node-link-protection (Protocols IS-IS)	287
nssa	288
options (Routing Options)	289
ospf	290
ospf3	291
out-delay	292
outbound-route-filter	293
overload (Protocols IS-IS)	295
overload (Protocols OSPF)	298
passive (Protocols BGP)	299
passive (Protocols IS-IS)	300
passive (Protocols OSPF)	301
peer-as (Protocols BGP)	303
pim-to-igmp-proxy	305
pim-to-mld-proxy	306
point-to-point	307
policy (Aggregate and Generated Routes)	308
policy (Flow Maps)	309
policy (SSM Maps)	310
ppm (Ethernet Switching)	311
ppm	312
preference (Protocols BGP)	313
preference (Protocols IS-IS)	314
preference (Protocols OSPF)	315
preference (Protocols RIP)	316
preference (Protocols RIPng)	317
preference (Routing Options)	318
prefix	319
prefix-export-limit (Protocols IS-IS)	320
prefix-export-limit (Protocols OSPF)	321
prefix-limit	322
priority (Protocols IS-IS)	324
priority (Protocols OSPF)	325
qualified-next-hop (Static Routes)	326
readvertise	328

realm	329
receive (Protocols RIP)	330
receive (Protocols RIPng)	331
redundant-sources	332
reference-bandwidth (Protocols IS-IS)	333
reference-bandwidth (Protocols OSPF)	334
register-probe-time	335
remove-private	336
resolution	338
resolution-ribs	339
resolve	340
restart-duration (Routing Options)	341
retain	342
retransmit-interval (OSPF)	343
reverse-oif-mapping	344
rib (General)	345
rib (Route Resolution)	347
rib-group (Protocols BGP)	348
rib-group (Protocols IS-IS)	349
rib-group (Protocols OSPF)	350
rib-group (Protocols RIP)	351
rib-group (Routing Options)	352
rib-groups	353
rip	355
ripng	355
route-distinguisher-id	356
route-record	357
route-timeout (Protocols RIP)	358
route-timeout (Protocols RIPng)	359
route-type-community	360
router-id	361
routing-options	362
rpf-check-policy (Routing Options RPF)	362
scope	363
scope-policy	364
send (Protocols RIP)	365
send (Protocols RIPng)	366
shortcuts (Protocols OSPF)	367
source (Source-Specific Multicast)	368
source-routing	369
spf-options (Protocols IS-IS)	370
spf-options (Protocols OSPF)	372
ssm-groups	374
ssm-map (Routing Options Multicast)	375
static (Routing Options)	376
stub	382
subscriber-leave-timer	383
summaries	384
tag (Routing Options)	385

	tcp-mss (Protocols BGP)	386
	threshold (Multicast Forwarding Cache)	387
	timeout (Flow Maps)	388
	timeout (Multicast)	389
	topologies (Protocols IS-IS)	390
	traceoptions (Protocols BGP)	391
	traceoptions (Protocols IS-IS)	394
	traceoptions (Protocols OSPF)	397
	traceoptions (Protocols RIP)	400
	traceoptions (Protocols RIPng)	403
	traceoptions (Routing Options)	406
	traffic-engineering (OSPF)	409
	transit-delay (OSPF)	411
	type (Protocols BGP)	412
	type-7	413
	update-interval (Protocols RIP)	414
	update-interval (Protocols RIPng)	415
	upstream-interface	416
	virtual-link	417
	wide-metrics-only	418
Chapter 5	Operational Commands for Layer 3 Protocols	419
	clear (ospf ospf3) database	420
	clear (ospf ospf3) io-statistics	423
	clear (ospf ospf3) neighbor	424
	clear (ospf ospf3) statistics	426
	clear bgp damping	428
	clear bgp neighbor	429
	clear bgp table	431
	clear ipv6 neighbors	433
	clear isis adjacency	434
	clear isis database	436
	clear isis overload	438
	clear isis statistics	440
	clear (ospf ospf3) overload	442
	clear rip general-statistics	443
	clear rip statistics	444
	clear ripng general-statistics	445
	clear ripng statistics	446
	show (ospf ospf3) interface	447
	show (ospf ospf3) io-statistics	453
	show (ospf ospf3) log	455
	show (ospf ospf3) neighbor	458
	show (ospf ospf3) overview	464
	show (ospf ospf3) route	469
	show (ospf ospf3) statistics	475
	show as-path	479
	show as-path domain	483
	show as-path summary	485

show bgp bmp	487
show bgp group	489
show bgp neighbor	496
show bgp summary	510
show ipv6 neighbors	515
show isis adjacency	517
show isis authentication	521
show isis backup coverage	523
show isis backup label-switched-path	525
show isis backup spf results	527
show isis database	531
show isis hostname	538
show isis interface	539
show isis overview	543
show isis route	546
show isis spf	550
show isis statistics	555
show ospf3 database	558
show ospf database	569
show policy damping	577
show rip general-statistics	579
show rip neighbor	581
show rip statistics	583
show ripng general-statistics	586
show ripng neighbor	587
show ripng statistics	589
show route	591
show route active-path	596
show route all	601
show route aspath-regex	603
show route best	605
show route brief	608
show route community	610
show route community-name	612
show route damping	614
show route detail	619
show route exact	634
show route export	636
show route extensive	639
show route flow validation	654
show route inactive-path	656
show route inactive-prefix	659
show route instance	661
show route label	668
show route label-switched-path	670
show route martians	672
show route next-hop	674
show route no-community	680
show route protocol	683

show route range	695
show route receive-protocol	699
show route resolution	707
show route snooping	710
show route source-gateway	718
show route summary	724
show route table	728
show route terse	739

List of Tables

Part 1	Layer 3 Protocols	
Chapter 1	Layer 3 Protocols—Overview	3
	Table 1: Supported Junos OS Layer 3 Protocol Statements and Features	3
	Table 2: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported	4
Chapter 2	Configuring Layer 3 Protocols	11
	Table 3: BGP Routing Configuration Summary	12
	Table 4: BGP Global Settings	14
	Table 5: OSPF Routing Configuration Summary	16
	Table 6: Edit OSPF Global Settings	18
	Table 7: RIP Routing Configuration Summary	20
	Table 8: Edit RIP Global Settings	21
	Table 9: Static Routing Configuration Summary	25
	Table 10: Policies Global Configuration Parameters	27
	Table 11: Terms Configuration Parameters	27
Chapter 3	Verifying Layer 3 Protocols Configuration	35
	Table 12: Summary of Key BGP Routing Output Fields	35
	Table 13: Summary of Key OSPF Routing Output Fields	38
	Table 14: Summary of Key RIP Routing Output Fields	40
	Table 15: Filtering Route Messages	42
	Table 16: Summary of Key Routing Information Output Fields	42
Chapter 4	Configuration Statements for Layer 3 Protocols	45
	Table 17: Unsupported [edit protocols bgp] Configuration Statements on EX Series Switches	54
	Table 18: Unsupported [edit protocols isis] Configuration Statements on EX Series Switches	58
	Table 19: Unsupported [edit protocols ospf] Configuration Statements on EX Series Switches	61
	Table 20: Unsupported [edit protocols ospf 3] Configuration Statements on EX Series Switches	64
	Table 21: Unsupported [edit protocols-rip] Configuration Statements on EX Series Switches	67
	Table 22: Default Metric Values for Routes Exported into IS-IS	247
Chapter 5	Operational Commands for Layer 3 Protocols	419
	Table 23: show (ospf ospf3) interface Output Fields	448
	Table 24: show (ospf ospf3) io-statistics Output Fields	453
	Table 25: show (ospf ospf3) log Output Fields	455

Table 26: show (ospf ospf3) neighbor Output Fields	459
Table 27: show ospf overview Output Fields	464
Table 28: show (ospf ospf3) route Output Fields	470
Table 29: show (ospf ospf3) statistics Output Fields	475
Table 30: show as-path Output Fields	479
Table 31: show as-path domain Output Fields	483
Table 32: show as-path summary Output Fields	485
Table 33: show bgp bmp Output Fields	487
Table 34: show bgp group Output Fields	490
Table 35: show bgp neighbor Output Fields	497
Table 36: show bgp summary Output Fields	510
Table 37: show ipv6 neighbors Output Fields	515
Table 38: show isis adjacency Output Fields	517
Table 39: show isis authentication Output Fields	521
Table 40: show isis backup coverage Output Fields	523
Table 41: show isis backup label-switched-path Output Fields	525
Table 42: show isis backup spf results Output Fields	528
Table 43: show isis database Output Fields	532
Table 44: show isis hostname Output Fields	538
Table 45: show isis interface Output Fields	540
Table 46: show isis overview Output Fields	543
Table 47: show isis route Output Fields	546
Table 48: show isis spf Output Fields	550
Table 49: show isis statistics Output Fields	556
Table 50: show ospf3 database Output Fields	559
Table 51: show ospf database Output Fields	570
Table 52: show policy damping Output Fields	577
Table 53: show rip general-statistics Output Fields	579
Table 54: show rip neighbor Output Fields	581
Table 55: show rip statistics Output Fields	584
Table 56: show ripng general-statistics Output Fields	586
Table 57: show ripng neighbor Output Fields	587
Table 58: show ripng statistics Output Fields	589
Table 59: show route Output Fields	592
Table 60: show route damping Output Fields	614
Table 61: show route detail Output Fields	619
Table 62: Next-hop Types Output Field Values	623
Table 63: State Output Field Values	625
Table 64: Communities Output Field Values	627
Table 65: show route export Output Fields	636
Table 66: show route extensive Output Fields	639
Table 67: show route flow validation Output Fields	654
Table 68: show route instance Output Fields	661
Table 69: show route martians Output Fields	672
Table 70: show route receive-protocol Output Fields	699
Table 71: show route resolution Output Fields	708
Table 72: show route summary Output Fields	724
Table 73: show route terse Output Fields	739

About This Topic Collection

- [How to Use This Guide on page xv](#)
- [List of EX Series Guides for Junos OS Release 12.3 on page xv](#)
- [Downloading Software on page xviii](#)
- [Documentation Symbols Key on page xviii](#)
- [Documentation Feedback on page xx](#)
- [Requesting Technical Support on page xx](#)

How to Use This Guide

Complete documentation for the EX Series product family is provided on webpages at http://www.juniper.net/techpubs/en_US/release-independent/information-products/pathway-pages/ex-series/product/index.html. We have selected content from these webpages and created a number of EX Series guides that collect related topics into a book-like format so that the information is easy to print and easy to download to your local computer.

Software features for EX Series switches are listed by platform and by Junos OS release in a standalone document. See [EX Series Switch Software Features Overview](#).

The release notes are at http://www.juniper.net/techpubs/en_US/junos12.1/information-products/topic-collections/release-notes/12.1/junos-release-notes-12.1.pdf.

List of EX Series Guides for Junos OS Release 12.3

Title	Description
<i>Complete Guide for the EX Series Redundant Power System</i>	Hardware descriptions, site preparation, installation, replacement, software features and configuration, and safety and compliance information for EX Series RPS.
<i>Complete Hardware Guide for EX2200 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX2200 Ethernet switches
<i>Complete Hardware Guide for EX3200 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX3200 Ethernet switches

Title	Description
<i>Complete Hardware Guide for EX3300 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX3300 Ethernet switches
<i>Complete Hardware Guide for EX4200 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX4200 Ethernet switches
<i>Complete Hardware Guide for EX4500 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX4500 Ethernet switches
<i>Complete Hardware Guide for EX6210 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX6210 Ethernet switches
<i>Complete Hardware Guide for EX8208 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8208 Ethernet switches
<i>Complete Hardware Guide for EX8216 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8216 Ethernet switches
<i>Complete Hardware Guide for the XRE200 External Routing Engine</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for the XRE200 External Routing Engine
<i>Complete Software Guide for Junos[®] OS for EX Series Ethernet Switches, Release 12.3</i>	Software feature descriptions, configuration examples, and tasks for Junos OS for EX Series switches
Software Topic Collections	Software feature descriptions, configuration examples and tasks, and reference pages for configuration statements and operational commands (This information also appears in the <i>Complete Software Guide for Junos[®] OS for EX Series Ethernet Switches, Release 12.3.</i>)
<hr/>	
<i>Junos[®] OS Access Privilege Configuration Guide</i>	
<i>Junos[®] OS for EX Series Ethernet Switches, Release 12.3: Access and User Management</i>	
<i>Junos[®] OS for EX Series Ethernet Switches, Release 12.3: Access Control</i>	
<i>Junos[®] OS for EX Series Ethernet Switches, Release 12.3: Configuration Management</i>	
<i>Junos[®] OS for EX Series Ethernet Switches, Release 12.3: Class of Service</i>	




Title	Description
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: Device Security</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: Ethernet Switching</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: EX3300, EX4200, and EX4500 Virtual Chassis</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: EX8200 Virtual Chassis</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: Converged Networks (LAN and SAN)</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: High Availability</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: Interfaces</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: Layer 3 Protocols</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: MPLS</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: Multicast</i>	
<i>Junos® OS for EX Series Switches, Release 12.3: Network Management and Monitoring</i>	
<i>Junos® OS for EX Series Switches, Release 12.3: Port Security</i>	
<i>Junos® OS for EX Series Switches, Release 12.3: Power over Ethernet</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: Routing Policy and Packet Filtering</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: Software Installation</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: Spanning-Tree Protocols</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: System Monitoring</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 12.3: System Services</i>	

Title	Description
<i>Junos[®] OS for EX Series Ethernet Switches, Release 12.3: System Setup</i>	
<i>Junos[®] OS for EX Series Ethernet Switches, Release 12.3: User Interfaces</i>	

Downloading Software

You can download Junos OS for EX Series switches from the Download Software area at <http://www.juniper.net/customers/support/>. To download the software, you must have a Juniper Networks user account. For information about obtaining an account, see <http://www.juniper.net/entitlement/setupAccountInfo.do>.

Documentation Symbols Key

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.

Text and Syntax Conventions		
Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active

Text and Syntax Conventions		
Convention	Description	Examples
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces important new terms. Identifies book names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [<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		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.

Text and Syntax Conventions

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

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. Send e-mail to techpubs-comments@juniper.net with the following:

- Document URL or title
- Page number if applicable
- Software version
- Your name and company

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

Layer 3 Protocols

- [Layer 3 Protocols—Overview on page 3](#)
- [Configuring Layer 3 Protocols on page 11](#)
- [Verifying Layer 3 Protocols Configuration on page 35](#)
- [Configuration Statements for Layer 3 Protocols on page 45](#)
- [Operational Commands for Layer 3 Protocols on page 419](#)

CHAPTER 1

Layer 3 Protocols—Overview

- [Layer 3 Protocols Supported on EX Series Switches](#) on page 3
- [Layer 3 Protocols Not Supported on EX Series Switches](#) on page 4
- [Understanding Distributed Periodic Packet Management on EX Series Switches](#) on page 6
- [Understanding IPsec Authentication for OSPF Packets on EX Series Switches](#) on page 7

Layer 3 Protocols Supported on EX Series Switches

EX Series switches support the Junos OS Layer 3 features and configuration statements listed in [Table 1 on page 3](#):

Table 1: Supported Junos OS Layer 3 Protocol Statements and Features

Protocol	Notes	For More Information
BGP	Fully supported.	Junos OS Routing Protocols Configuration Guide
BFD	Fully supported.	Junos OS Routing Protocols Configuration Guide
ICMP	Fully supported.	Junos OS Routing Protocols Configuration Guide
IGMPv1, v2, and v3	Fully supported.	Junos OS Multicast Protocols Configuration Guide
IS-IS	Supported, with the exceptions noted in “ Layer 3 Protocols Not Supported on EX Series Switches ” on page 4.	Junos OS Routing Protocols Configuration Guide
MLD	Fully supported (MLD versions 1 and 2).	Junos OS Multicast Protocols Configuration Guide
MPLS	Supported, with the exceptions noted in “ Layer 3 Protocols Not Supported on EX Series Switches ” on page 4.	Junos OS MPLS Applications Configuration Guide
OSPFv1, v2 and v3	Supported, with the exceptions noted in “ Layer 3 Protocols Not Supported on EX Series Switches ” on page 4.	Junos OS Routing Protocols Configuration Guide

Table 1: Supported Junos OS Layer 3 Protocol Statements and Features (*continued*)

Protocol	Notes	For More Information
PIM	Fully supported on EX3200, EX3300, EX4200, EX6200, and EX8200 switches.	Junos OS Multicast Protocols Configuration Guide
PPM	Supported. See EX Series Switch Software Features Overview for specific platform information.	Junos OS Routing Protocols Configuration Guide
RIP	Fully supported.	Junos OS Routing Protocols Configuration Guide
RIPng	Fully supported.	Junos OS Routing Protocols Configuration Guide
SNMP	Fully supported.	Junos OS Network Management Configuration Guide
VRRP	Fully supported.	See Understanding VRRP on EX Series Switches. See also Junos OS High Availability Guide .

- Related Documentation**
- [Layer 3 Protocols Not Supported on EX Series Switches on page 4](#)
 - EX Series Switch Software Features Overview

Layer 3 Protocols Not Supported on EX Series Switches

EX Series switches do not support the Junos OS Layer 3 protocols and features listed in [Table 2 on page 4](#):

Table 2: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported

Feature	Configuration Statements Not Supported on EX Series Switches
DVMRP	<ul style="list-style-type: none"> • dvmrp and subordinate statements
Flow aggregation (cflowd)	<ul style="list-style-type: none"> • cflow and subordinate statements
IPsec	<ul style="list-style-type: none"> • [edit services] statements related to IPsec
IS-IS: <ul style="list-style-type: none"> • ES-IS • IPv6 in multicast routing protocols 	<ul style="list-style-type: none"> • clns-routing statement • ipv6-multicast statement • lsp-interval statement • label-switched-path statement • lsp-lifetime statement • te-metric statement
Logical routers	<ul style="list-style-type: none"> • logical-routers and subordinate statements

Table 2: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)

Feature	Configuration Statements Not Supported on EX Series Switches
MPLS: <ul style="list-style-type: none"> Fast Reroute (FRR) Label Distribution Protocol (LDP) (except on EX8200 switches) Layer 3 VPNs (except on EX8200 switches) Multiprotocol BGP (MP-BGP) for VPN-IPv4 family Pseudowire emulation (PWE3) Routing policy statements related to Layer 3 VPNs and MPLS (except on EX8200 switches) Virtual Private LAN Service (VPLS) 	<ul style="list-style-type: none"> ldp and all subordinate statements (except on EX8200 switches)
Network Address Translation (NAT)	<ul style="list-style-type: none"> nat and subordinate statements Policy statements related to NAT
OSPF	<ul style="list-style-type: none"> demand-circuit statement label-switched-path and subordinate statements neighbor statement within an OSPF area peer-interface and subordinate statements within an OSPF area sham-link statement te-metric statement
PIM SM	<ul style="list-style-type: none"> Not supported on EX2200 switches
PIM SSM	<ul style="list-style-type: none"> Not supported on EX2200 switches
PIM DM	<ul style="list-style-type: none"> Not supported on EX2200 or EX4500 switches
PIM: <ul style="list-style-type: none"> IPv6 	<ul style="list-style-type: none"> inet6 family (EX2200 and EX4500 switches)
PPM	<ul style="list-style-type: none"> Not supported on EX2200 and EX3300 switches
Routing instances: <ul style="list-style-type: none"> Routing instance forwarding 	<ul style="list-style-type: none"> l2vpn and subordinate statements (except on EX8200 switches) ldp and subordinate statements (except on EX8200 switches) vpls and subordinate statements
Routed VLAN interfaces (RVIs)	<ul style="list-style-type: none"> family mpls statement

Table 2: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)

Feature	Configuration Statements Not Supported on EX Series Switches
SAP and SDP	<ul style="list-style-type: none"> • sap and all subordinate statements
General routing options in the routing-options hierarchy: <ul style="list-style-type: none"> • MPLS and label-switched-paths 	<ul style="list-style-type: none"> • auto-export and subordinate statements • dynamic-tunnels and subordinate statements • lsp-next-hop and subordinate statements • multicast and subordinate statements • p2mp-lsp-next-hop and subordinate statements • route-distinguisher-id statement (except on EX8200 switches)
Traffic sampling and forwarding in the forwarding-options hierarchy	<ul style="list-style-type: none"> • accounting and subordinate statements • family mpls and family multiservice under hash-key hierarchy • Under monitoring group-name family inet output hierarchy: <ul style="list-style-type: none"> • cflowd statement • export-format-cflowd-version-5 statement • flow-active-timeout statement • flow-export-destination statement • flow-inactive-timeout statement • interface statement • port-mirroring statement (On EX Series switches, port mirroring is implemented using the analyzer statement.) • sampling and subordinate statements

- Related Documentation**
- [Layer 3 Protocols Supported on EX Series Switches on page 3](#)
 - [EX Series Switch Software Features Overview](#)

Understanding Distributed Periodic Packet Management on EX Series Switches

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks for particular processes so that other processes on the Juniper Networks EX Series Ethernet Switch can more optimally direct their resources. PPM is responsible for the periodic transmission of packets on behalf of its various client processes, which include the processes that control the Link Aggregation Control Protocol (LACP) and Bidirectional Forwarding Detection (BFD) protocols, and also for receiving packets on behalf of these client processes. PPM also gathers some statistics and sends process-specific packets. PPM cannot be disabled and is always running on any operational switch.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and either the access interfaces (on EX3200, EX4200, and EX4500 switches) or the line cards (on EX6200 and EX8200 switches) for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for LACP packets only.



BEST PRACTICE: We recommend that, generally, you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

- Related Documentation**
- [Configuring Distributed Periodic Packet Management on an EX Series Switch \(CLI Procedure\) on page 31](#)

Understanding IPsec Authentication for OSPF Packets on EX Series Switches

IP Security (IPsec) provides a secure way to authenticate senders and encrypt IP version 4 (IPv4) traffic between network devices. IPsec offers network administrators for Juniper Networks EX Series Ethernet Switches and their users the benefits of data confidentiality, data integrity, sender authentication, and anti-replay services.

IPsec is a framework for ensuring secure private communication over IP networks and is based on standards developed by the International Engineering Task Force (IETF). IPsec provides security services at the network layer of the Open Systems Interconnection (OSI) model by enabling a system to select required security protocols, determine the algorithms to use for the security services, and implement any cryptographic keys required to provide the requested services. You can use IPsec to protect one or more paths between a pair of hosts, between a pair of security gateways (such as switches), or between a security gateway and a host.

OSPF version 3 (OSPFv3), unlike OSPF version 2 (OSPFv2), does not have a built-in authentication method and relies on IPsec to provide this functionality. You can secure specific OSPFv3 interfaces and protect OSPFv3 virtual links.

- [Authentication Algorithms on page 7](#)
- [Encryption Algorithms on page 8](#)
- [IPsec Protocols on page 8](#)
- [Security Associations on page 9](#)
- [IPsec Modes on page 9](#)

Authentication Algorithms

Authentication is the process of verifying the identity of the sender. Authentication algorithms use a shared key to verify the authenticity of the IPsec devices. The Juniper Networks Junos operating system (Junos OS) uses the following authentication algorithms:

- Message Digest 5 (MD5) uses a one-way hash function to convert a message of arbitrary length to a fixed-length message digest of 128 bits. Because of the conversion process, it is mathematically infeasible to calculate the original message by computing it

backwards from the resulting message digest. Likewise, a change to a single character in the message will cause it to generate a very different message digest number.

To verify that the message has not been tampered with, Junos OS compares the calculated message digest against a message digest that is decrypted with a shared key. Junos OS uses the MD5 hashed message authentication code (HMAC) variant that provides an additional level of hashing. MD5 can be used with an authentication header (AH) and Encapsulating Security Payload (ESP).

- Secure Hash Algorithm 1 (SHA-1) uses a stronger algorithm than MD5. SHA-1 takes a message of less than 264 bits in length and produces a 160-bit message digest. The large message digest ensures that the data has not been changed and that it originates from the correct source. Junos OS uses the SHA-1 HMAC variant that provides an additional level of hashing. SHA-1 can be used with AH, ESP, and Internet Key Exchange (IKE).

Encryption Algorithms

Encryption encodes data into a secure format so that it cannot be deciphered by unauthorized users. As with authentication algorithms, a shared key is used with encryption algorithms to verify the authenticity of IPsec devices. Junos OS uses the following encryption algorithms:

- Data Encryption Standard cipher-block chaining (DES-CBC) is a symmetric secret-key block algorithm. DES uses a key size of 64 bits, where 8 bits are used for error detection and the remaining 56 bits provide encryption. DES performs a series of simple logical operations on the shared key, including permutations and substitutions. CBC takes the first block of 64 bits of output from DES, combines this block with the second block, feeds this back into the DES algorithm, and repeats this process for all subsequent blocks.
- Triple DES-CBC (3DES-CBC) is an encryption algorithm that is similar to DES-CBC but provides a much stronger encryption result because it uses three keys for 168-bit (3 x 56-bit) encryption. 3DES works by using the first key to encrypt the blocks, the second key to decrypt the blocks, and the third key to reencrypt the blocks.

IPsec Protocols

IPsec protocols determine the type of authentication and encryption applied to packets that are secured by the switch. Junos OS supports the following IPsec protocols:

- AH—Defined in *RFC 2402*, AH provides connectionless integrity and data origin authentication for IPv4. It also provides protection against replays. AH authenticates as much of the IP header as possible, as well as the upper-level protocol data. However, some IP header fields might change in transit. Because the value of these fields might not be predictable by the sender, they cannot be protected by AH. In an IP header, AH can be identified with a value of 51 in the Protocol field of an IPv4 packet.
- ESP—Defined in *RFC 2406*, ESP can provide encryption and limited traffic flow confidentiality or connectionless integrity, data origin authentication, and an anti-replay service. In an IP header, ESP can be identified with a value of 50 in the Protocol field of an IPv4 packet.

Security Associations

An IPsec consideration is the type of security association (SA) that you wish to implement. An SA is a set of IPsec specifications that are negotiated between devices that are establishing an IPsec relationship. These specifications include preferences for the type of authentication, encryption, and IPsec protocol to be used when establishing the IPsec connection. An SA can be either unidirectional or bidirectional, depending on the choices made by the network administrator. An SA is uniquely identified by a Security Parameter Index (SPI), an IPv4 or IPv6 destination address, and a security protocol (AH or ESP) identifier.

IPsec Modes

Junos OS supports the following IPsec modes:

- Tunnel mode is supported for both AH and ESP in Junos OS. In tunnel mode, the SA and associated protocols are applied to tunneled IPv4 or IPv6 packets. For a tunnel mode SA, an outer IP header specifies the IPsec processing destination and an inner IP header specifies the ultimate destination for the packet. The security protocol header appears after the outer IP header and before the inner IP header. In addition, there are slight differences for tunnel mode when you implement it with AH and ESP:
 - For AH, portions of the outer IP header are protected, as well as the entire tunneled IP packet.
 - For ESP, only the tunneled packet is protected, not the outer header.

When one side of an SA is a security gateway (such as a switch), the SA must use tunnel mode. However, when traffic (for example, SNMP commands or BGP sessions) is destined for a switch, the system acts as a host. Transport mode is allowed in this case because the system does not act as a security gateway and does not send or receive transit traffic.



NOTE: Tunnel mode is not supported for OSPF v3 control packet authentication.

- Transport mode provides an SA between two hosts. In transport mode, the protocols provide protection primarily for upper-layer protocols. A transport mode security protocol header appears immediately after the IP header and any options and before any higher-layer protocols (for example, TCP or UDP). There are slight differences for transport mode when you implement it with AH and ESP:
 - For AH, selected portions of the IP header are protected, as well as selected portions of the extension headers and selected options within the IPv4 header.
 - For ESP, only the higher-layer protocols are protected, not the IP header or any extension headers preceding the ESP header.

Related Documentation

- [Using IPsec to Secure OSPFv3 Networks \(CLI Procedure\) on page 32](#)
- [Configuring an OSPF Network \(J-Web Procedure\) on page 15](#)

CHAPTER 2

Configuring Layer 3 Protocols

- [Configuring BGP Sessions \(J-Web Procedure\) on page 11](#)
- [Configuring an OSPF Network \(J-Web Procedure\) on page 15](#)
- [Configuring a RIP Network \(J-Web Procedure\) on page 19](#)
- [Configuring Static Routing \(CLI Procedure\) on page 24](#)
- [Configuring Static Routing \(J-Web Procedure\) on page 24](#)
- [Configuring Routing Policies \(J-Web Procedure\) on page 26](#)
- [Configuring Distributed Periodic Packet Management on an EX Series Switch \(CLI Procedure\) on page 31](#)
- [Using IPsec to Secure OSPFv3 Networks \(CLI Procedure\) on page 32](#)

Configuring BGP Sessions (J-Web Procedure)

You can use the J-Web interface to create BGP peering sessions on a routing device.



NOTE: To configure BGP sessions, you must have a license for BGP installed on the EX Series switch.

To configure a BGP peering session:

1. Select **Configure > Routing > BGP**.



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

2. Click one:

- **Add**—Adds a BGP group. Enter information into the configuration page as described in [Table 3 on page 12](#).
- **Edit**—Modifies an existing BGP group. Enter information into the configuration page as described in [Table 3 on page 12](#).
- **Delete**—Deletes an existing BGP group.
- **Disable**—Disables BGP configuration.

3. To modify BGP global settings, click **Edit** in the Global Information section. Enter information as described in [Table 4 on page 14](#).**Table 3: BGP Routing Configuration Summary**

Field	Function	Your Action
General tab		
Group Type	Specifies whether the group is an internal BGP (IBGP) group or an external BGP (EBGP) group.	Select the option: Internal or External .
Group Name	Specifies the name for the group.	Type a new name or select and edit the name.
ASN	Sets the unique numeric identifier of the AS in which the routing device is configured.	Type the routing device's 32-bit AS number, in dotted decimal notation. If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3 , the value assigned to the AS is 0.0.0.3 .
Preference	Specifies the degree of preference for an external route. The route with the highest local preference value is preferred.	Type or select and edit the value.
Cluster Id	Specifies the cluster identifier to be used by the route reflector cluster in an internal BGP group.	Type or select and edit the IPv6 or IPv4 address to be used as the identifier.
Description	Specifies the text description of the global, group, or neighbor configuration.	Type or select and edit the description.
Damping	Specifies whether route flap damping is enabled or not.	To enable route flap damping, select the check box. To disable route flap damping do not select the check box.
Advertise Inactive Routes	Specifies whether BGP advertises the best route even if the routing table did not select it to be an active route.	To enable advertising inactive routes, select the check box. To disable advertising inactive routes, do not select the check box.

Table 3: BGP Routing Configuration Summary (*continued*)

Field	Function	Your Action
Advertise Peer AS Routes	Specifies whether to disable the default behavior of suppressing AS routes.	To enable advertising peer AS routes, select the check box. To disable advertising peer AS routes, do not select the check box.
Neighbors tab		
Dynamic Neighbors	Configures a neighbor (peer).	Type the IPv4 address of the peer.
Static Neighbors	Configures the system's peers statically.	<p>To configure a static neighbor:</p> <ol style="list-style-type: none"> 1. Specify the IP address. 2. Specify the address of the local end of a BGP session. 3. Specify the degree of preference for an external route. 4. Enter a description. 5. Specify the hold-time value to use when negotiating a connection with the peer. 6. Specify how long a route must be present in the routing table before it is exported to BGP. Use this time delay to help bundle routing updates. 7. Select Passive if you do not want to send active open messages to the peer. 8. Select the option to compare the AS path of an incoming advertised route with the AS number of the BGP peer under the group and replace all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer. 9. Specify an import policy and export policy. 10. Click OK.
Policies tab		
Import Policy	Specifies one or more routing policies to routes being imported into the routing table from BGP.	<p>Click Add to add an import policy. Select the policy and click OK.</p> <p>Click Move up or Move down to move the selected policy up or down the list of policies.</p> <p>Select the policy and click Remove.</p>
Export Policy	Specifies one or more policies to routes being exported from the routing table into BGP.	<p>Click Add to add an export policy. Select the policy and click OK.</p> <p>Click Move up or Move down to move the selected policy up or down the list of policies.</p> <p>Select the policy and click Remove.</p>

Table 4: BGP Global Settings

Field	Function	Your Action
General tab		
Router ASN	Specifies the routing device's AS number.	Type or select and edit the value.
Router Identifier	Specify the routing device's IP address.	Type or select and edit the IP address.
BGP Status	Enables or disables BGP.	<ul style="list-style-type: none"> To enable BGP, select Enabled. To disable BGP, select Disabled.
Description	Describes of the global, group, or neighbor configuration.	Type or select and edit the description.
Confederation Number	Specifies the routing device's confederation AS number.	Type or select and edit the value.
Confederation Members	Specifies the AS numbers for the confederation members.	<p>To add a member AS number, click Add and enter the number in the Member ASN box. Click OK.</p> <p>To modify a confederation member's AS number, select the member click Edit and, enter the number and click OK.</p> <p>To delete a confederation member, select the member and click Remove.</p>
Advance Options	<p>You can configure the following:</p> <ul style="list-style-type: none"> Keep routes—Specifies whether routes learned from a BGP peer must be retained in the routing table even if they contain an AS number that was exported from the local AS. TCP MSS—Configures the maximum segment size (MSS) for the TCP connection for BGP neighbors. MTU Discovery—Select to configure MTU discovery. Remove Private ASN—Select to have the local system strip private AS numbers from the AS path when advertising AS paths to remote systems. Graceful Restart—Specifies the time period when the restart is expected to be complete. Specify the maximum time that stale routes are kept during restart. Multihop—Configures the maximum time-to-live (TTL) value for the TTL in the IP header of BGP packets. Authentication Type—Select the authentication algorithm: None, MD5, SHA1, AES. 	<p>Select All or None to configure Keep Routes.</p> <p>Enter a value in the TCP MSS box.</p> <p>Click to enable MTU Discovery.</p> <p>Click to enable Remove Private ASN.</p> <p>Enter the time period for a graceful restart and the maximum time that stale routes must be kept.</p> <p>To configure Multihop, select Nexthop Change to allow unconnected third-party next hops. Enter a TTL value.</p> <p>Select the authentication algorithm. If you select None, specify an authentication key (password).</p>
Policies tab		

Table 4: BGP Global Settings (*continued*)

Field	Function	Your Action
Import Policy	Specifies one or more routing policies to routes being imported into the routing table from BGP.	<p>Click Add to add an import policy.</p> <p>Click Move up or Move down to move the selected policy up or down the list of policies.</p> <p>Click Remove to remove an import policy.</p>
Export Policy	Specifies one or more policies to routes being exported from the routing table into BGP.	<p>Click Add to add an export policy.</p> <p>Click Move up or Move down to move the selected policy up or down the list of policies.</p> <p>Click Remove to remove an export policy.</p>
Trace Options tab		
File Name	Specifies the name of the file to receive the output of the tracing operation.	Type or select and edit the name.
Number of Files	Specifies the maximum number of trace files.	Type or select and edit the value.
File Size	Specifies the maximum size for each trace file.	Type or select and edit the value.
World Readable	Specifies whether the trace file can be read by any user or not.	<p>Select True to allow any user to read the file.</p> <p>Select False to disallow all users being able to read the file.</p>
Flags	Specifies the tracing operation to perform.	Select a value from the list.

- Related Documentation**
- [Monitoring BGP Routing Information on page 35](#)
 - [Layer 3 Protocols Supported on EX Series Switches on page 3](#)

Configuring an OSPF Network (J-Web Procedure)

You can use the J-Web interface to create multiarea OSPF networks on an EX Series switch.

To configure a multiarea OSPF network:

1. Select **Configure > Routing > OSPF**.



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

2. Click one:
 - **Add**—Adds an OSPF area. Enter information into the configuration page as described in [Table 5 on page 16](#).
 - **Edit**—Modifies an existing OSPF area. Enter information into the configuration page as described in [Table 5 on page 16](#).
 - **Delete**—Deletes an existing OSPF area.
3. To modify OSPF global settings, click **Edit**. Enter information as described in [Table 6 on page 18](#).
4. To disable OSPF, click **Disable**.

Table 5: OSPF Routing Configuration Summary

Field	Function	Your Action
General tab		
Area Id	Uniquely identifies the area within its AS.	<p>Type a 32-bit numeric identifier for the area. Type an integer or select and edit the value.</p> <p>If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the area is 0.0.0.3.</p>
Area Ranges	Specifies a range of IP addresses for an area when sending summary link advertisements (within an area).	<p>To add a range:</p> <ol style="list-style-type: none"> 1. Click Add. 2. Type the area range. 3. Specify the subnet mask. 4. To override the metric for the IP address range, type a specific metric value. 5. If you do not want to display the routes that are contained within a summary, select Restrict advertisements of this area range. 6. If you want a summary of a route to be advertised only when an exact match is made with the configured summary range, select Enforce exact match for advertisement of this area range. 7. Click OK. <p>To modify an existing area range, select the area range, click Edit, and edit the value. Click OK.</p> <p>To delete an area range, select the area range and click Delete.</p>

Table 5: OSPF Routing Configuration Summary (*continued*)

Field	Function	Your Action
Area Type	<p>Designates the type of OSPF area.</p> <ul style="list-style-type: none"> • regular—A regular OSPF area, including the backbone area • stub—A stub area • nssa—A not-so-stubby area (NSSA) 	<p>Select the type of OSPF area you are creating from the list.</p> <p>If you select stub:</p> <ol style="list-style-type: none"> 1. Enter the default metric. 2. To flood summary LSAs into the stub area, select the check box. <p>If you select nssa:</p> <ol style="list-style-type: none"> 1. Specify the metric type. 2. Enter the default metric. 3. To flood summary LSAs into the nssa area, select the check box. 4. To flood Type-7 LSAs into the nssa area, select the check box.
Interfaces tab		
Interfaces	Specifies the interfaces to be associated with the OSPF configuration	<p>To associate an interface with the configuration, select the interface from the list, select Associate and click OK.</p> <p>To edit an interface's configuration:</p> <ol style="list-style-type: none"> 1. Select the interface from the list and click Edit. 2. Specify the cost of an OSPF interface. 3. Specify the traffic engineering metric. 4. Specify how often the routing device sends hello packets from the interface. 5. Specify how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements to an interface's neighbors. 6. To enable OSPF on the interface, select the check box. 7. To inform other protocols about neighbor down events, select the check box. 8. To treat the interface as a secondary interface, select the check box. 9. To only advertise OSPF, select the check box. 10. Click OK.
Policies tab		
Import Policy	Specifies one or more policies to control which routes learned from an area are used to generate summary link-state advertisements (LSAs) into other areas.	<p>Click Add to add an import policy.</p> <p>Click Move up or Move down to move the selected policy up or down the list of policies.</p> <p>Click Remove to remove an import policy.</p>

Table 5: OSPF Routing Configuration Summary (*continued*)

Field	Function	Your Action
Export Policy	Specifies one or more policies to control which summary LSAs are flooded into an area.	<p>Click Add to add an export policy.</p> <p>Click Move up or Move down to move the selected policy up or down the list of policies.</p> <p>Click Remove to remove an export policy.</p>

Table 6: Edit OSPF Global Settings

Field	Function	Your Action
General tab		
Router Id	Specifies the ID for the routing device.	Type or select and edit the value.
RIB Group	Installs the routes learned from OSPF routing instances into routing tables in the OSPF routing table group.	Select a value.
Internal Route Preference	Specifies the route preference for internal groups.	Type or select and edit the value.
External Route Preference	Specifies the route preference for external groups.	Type or select and edit the value.
Graceful Restart	Configures graceful restart for OSPF.	<p>To configure graceful restart:</p> <ol style="list-style-type: none"> 1. Specify the estimated time to send out purged grace LSAs over all the interfaces. 2. Specified the estimated time to reacquire a full OSPF neighbor from each area. 3. To disable No Strict LSA Checking, select the check box. 4. To disable graceful restart helper capability, select the check box. Helper mode is enabled by default. 5. Click OK.
SPF Options	Configure options for running the shortest-path-first (SPF) algorithm. You can configure a delay for when to run the SPF algorithm after a network topology change is detected, the maximum number of times the SPF algorithm can run in succession, and a hold-down interval after the SPF algorithm runs the maximum number of times.	<p>To configure SPF:</p> <ol style="list-style-type: none"> 1. Specify the time interval between the detection of a topology change and when the SPF algorithm runs. 2. Specify the time interval to hold down, or wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession. 3. Specify the maximum number of times the SPF algorithm can run in succession. After the maximum is reached, the hold-down interval begins.

Table 6: Edit OSPF Global Settings (*continued*)

Field	Function	Your Action
Policies tab		
Import Policy	Specifies one or more policies to control which routes learned from an area are used to generate summary link-state advertisements (LSAs) into other areas.	<p>Click Add to add an import policy.</p> <p>Click Move up or Move down to move the selected policy up or down the list of policies.</p> <p>Click Remove to remove an import policy.</p>
Export Policy	Specifies one or more policies to control which summary LSAs are flooded into an area.	<p>Click Add to add an export policy.</p> <p>Click Move up or Move down to move the selected policy up or down the list of policies.</p> <p>Click Remove to remove an export policy.</p>
Trace Options tab		
File Name	Specifies the name of the file to receive the output of the tracing operation.	Type or select and edit the name.
Number of Files	Specifies the maximum number of trace files.	Type or select and edit the name.
File Size	Specifies the maximum size for each trace file.	Type or select and edit the name.
World Readable	Specifies whether the trace file can be read by any user or not.	<p>Select True to allow any user to read the file.</p> <p>Select False to disallow all users being able to read the file.</p>
Flags	Specifies the tracing operation to perform.	Select a value from the list.

- Related Documentation**
- [Monitoring OSPF Routing Information on page 37](#)
 - [Layer 3 Protocols Supported on EX Series Switches on page 3](#)

Configuring a RIP Network (J-Web Procedure)

You can use the J-Web interface to create RIP networks.

To configure a RIP network:

1. Select **Configure > Routing > RIP**.



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

2. Click one:
 - **Add**—Configures a RIP instance. Enter information into the RIP Configuration page as described in [Table 7 on page 20](#).
 - **Edit**—Modifies an existing RIP instance. Enter information into the configuration page for RIP as described in [Table 7 on page 20](#).
 - **Delete**—Deletes an existing RIP instance.
4. To modify RIP global settings, click **Edit**. Enter information in the configuration as described in [Table 8 on page 21](#).

Table 7: RIP Routing Configuration Summary

Field	Function	Your Action
General tab		
Routing instance name	Specifies a name for the routing instance.	Type or select and edit the name.
Preference	Specifies the preference of external routes learned by RIP as compared to those learned from other routing protocols.	Type or select and edit the value.
Metric Out	Specifies the metric value to add to routes transmitted to the neighbor.	Type or select and edit the value.
Update interval	Specifies an update time interval to periodically send out routes learned by RIP to neighbors.	Type or select and edit the value.
Route timeout	Specifies the route timeout interval for RIP.	Type or select and edit the value.
Policies tab		
Import Policy	Applies one or more policies to routes being imported into the local routing device from the neighbors.	<p>Click Add to add an import policy.</p> <p>Click Move up or Move down to move the selected policy up or down the list of policies.</p> <p>Click Remove to remove an import policy.</p>

Table 7: RIP Routing Configuration Summary (*continued*)

Field	Function	Your Action
Export Policy	Applies a policy to routes being exported to the neighbors.	Click Add to add an export policy. Click Move up or Move down to move the selected policy up or down the list of policies. Click Remove to remove an export policy.
Neighbors tab		
RIP-Enabled Interfaces	Selects the interfaces to be associated with the RIP instance.	To enable RIP on an interface, click the check box next to the interface name. Click Edit if you want to modify an interface's settings.

Table 8: Edit RIP Global Settings

Field	Function	Your Action
General tab		
Send	Specifies RIP send options.	Select a value.
Receive	Configure RIP receive options.	Select a value.
Route timeout (sec)	Specifies the route timeout interval for RIP.	Type a value.
Update interval (sec)	Specifies the update time interval to periodically send out routes learned by RIP to neighbors.	Type or select and edit the value.
Hold timeout (sec)	Specifies the time period the expired route is retained in the routing table before being removed.	Type or select and edit the value.
Metric in	Specifies the metric to add to incoming routes when advertising into RIP routes that were learned from other protocols.	Type or select and edit the value.
RIB Group	Specifies a routing table group to install RIP routes into multiple routing tables.	Select and edit the name of the routing table group.
Message size	Specifies the number of route entries to be included in every RIP update message.	Type or select and edit the value.

Table 8: Edit RIP Global Settings (*continued*)

Field	Function	Your Action
Check Zero	<p>Specifies whether the reserved fields in a RIP packet are zero. Options are:</p> <ul style="list-style-type: none"> • check-zero—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications. • no-check-zero—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453. 	Select a value.
Graceful switchover	Configures graceful switchover for OSPF.	<p>To disable graceful restart, select Disable.</p> <p>Type or select and edit the estimated time for the restart to finish, in seconds.</p>
Authentication Type	<p>Specifies the type of authentication for RIP route queries received on an interface. Options are:</p> <ul style="list-style-type: none"> • None • MD5 • Simple 	<p>Select the authentication type.</p> <p>Enter the authentication key for MD5.</p>
Policies tab		
Import Policy	Applies one or more policies to routes being imported into the local routing device from the neighbors.	<p>Click Add to add an import policy.</p> <p>Click Move up or Move down to move the selected policy up or down the list of policies.</p> <p>Click Remove to remove an import policy.</p>
Export Policy	Applies a policy to routes being exported to the neighbors.	<p>Click Add to add an export policy.</p> <p>Click Move up or Move down to move the selected policy up or down the list of policies.</p> <p>Click Remove to remove an export policy.</p>
Trace Options tab		
File Name	Specifies the name of the file to receive the output of the tracing operation.	Type or select and edit the name.
Number of Files	Specifies the maximum number of trace files.	Type or select and edit the name.
File Size	Specifies the maximum size for each trace file.	Type or select and edit the name.

Table 8: Edit RIP Global Settings (*continued*)

Field	Function	Your Action
World Readable	Specifies whether the trace file can be read by any user or not.	Select True to allow any user to read the file. Select False to disallow all users being able to read the file.
Flags	Specifies the tracing operation to perform.	Select a value from the list.

- Related Documentation**
- [Monitoring RIP Routing Information on page 40](#)
 - [Layer 3 Protocols Supported on EX Series Switches on page 3](#)

Configuring Static Routing (CLI Procedure)

Static routes are routes that are manually configured and entered into the routing table. Dynamic routes, in contrast, are learned by the EX Series switch and added to the routing table using a protocol such as OSPF or RIP.

The switch uses static routes:

- When the switch does not have a route to a destination that has a better (lower) *preference* value. The preference is an arbitrary value in the range from 0 through 255 that the software uses to rank routes received from different protocols, interfaces, or remote systems. The routing protocol process generally determines the active route by selecting the route with the lowest preference value. In the given range, **0** is the lowest and **255** is the highest.
- When the switch cannot determine the route to a destination.
- When the switch is forwarding unroutable packets.

To configure basic static route options using the CLI:

- To configure the switch's default gateway:

```
[edit]
user@switch# set routing-options static route 0.0.0.0/0 next-hop 10.0.1.1
```

- To configure a static route and specify the next address to be used when routing traffic to the static route:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 next-hop 10.0.0.2.1
```

- To always keep the static route in the forwarding table:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 retain
```

- To prevent the static route from being readvertised:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 no-readvertise
```

- To remove inactive routes from the forwarding table:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 active
```

Related Documentation

- [Configuring Static Routing \(J-Web Procedure\) on page 24](#)
- [Monitoring Routing Information on page 41](#)

Configuring Static Routing (J-Web Procedure)

You can use the J-Web interface to configure static routes for EX Series switches.

To configure static routes:

1. Select **Configure > Routing > Static Routing**. The Static Routing page displays details of the configured routes.



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

2. Click one:
 - **Add**—To configure a route. Enter information into the routing page as described in [Table 9 on page 25](#).
 - **Edit**—To modify an existing route. Enter information into the routing page as described in [Table 9 on page 25](#).
 - **Delete**—To delete an existing route.

Table 9: Static Routing Configuration Summary

Field	Function	Your Action
Default Route		
Default Route	Specifies the default gateway for the switch. NOTE: IPv6 is not supported on EX2200 and EX4500 switches.	To specify an IPv4 address: <ol style="list-style-type: none"> 1. Select IPv4. 2. Type an IP address—for example, 10.10.10.10. 3. Enter the subnet mask or address prefix. For example, 24 bits represents 255.255.255.0. To specify an IPv6 address: <ol style="list-style-type: none"> 1. Select IPv6. 2. Type an IP address—for example, 2001:ab8:85a3::8a2e:370:7334. 3. Enter the subnet mask or address prefix.
Static Routes		

Table 9: Static Routing Configuration Summary (*continued*)

Field	Function	Your Action
Nexthop	Specifies the next-hop address or addresses to be used when routing traffic to the static route.	<p>To add an address:</p> <ol style="list-style-type: none"> 1. Click Add. 2. In the IP address dialog, enter the IP address. <p>NOTE: If a route has multiple next-hop addresses, traffic is routed across each address in round-robin fashion.</p> <ol style="list-style-type: none"> 3. Click OK. <p>To delete a next-hop address, select it from the list and click Delete.</p>

Related Documentation

- [Configuring Static Routing \(CLI Procedure\) on page 24](#)
- [Monitoring Routing Information on page 41](#)
- [Layer 3 Protocols Supported on EX Series Switches on page 3](#)

Configuring Routing Policies (J-Web Procedure)

All routing protocols use the Junos OS routing table to store the routes that they learn and to determine which routes are advertised in the protocol packets. Routing policy allows you to control which routes the routing protocols store in and retrieve from the routing table on the routing device.

To configure routing policies for an EX Series switch using the J-Web interface:

1. Select **Configure > Routing > Policies**.



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

2. Click one:

- **Global Options**—Configures global options for policies. Enter information into the configuration page as described in [Table 10 on page 27](#).
- **Add**—Configures a new policy. Select **New** and specify a policy name. To add terms, enter information into the configuration page as described in [Table 11 on page 27](#). Select **Clone** to create a copy of an existing policy.
- **Edit**—Edits an existing policy. To modify an existing term, enter information into the configuration page as described in [Table 11 on page 27](#).

- **Term Up**—Moves a term up in the list.
- **Term Down**—Moves a term down in the list.
- **Delete**—Deletes the selected policy.
- **Test Policy**—Tests the policy. Use this option to check whether the policy produces the results that you expect.

Table 10: Policies Global Configuration Parameters

Field	Function	Your Action
Prefix List	Specifies a list of IPv4 address prefixes for use in a routing policy statement.	<p>To add a prefix list:</p> <ol style="list-style-type: none"> 1. Click Add. 2. Enter a name for the prefix list. 3. To add an IP address, click Add. 4. Enter the IP address and the subnet mask and click OK. 5. Click OK. <p>To edit a prefix list, click Edit. Edit the settings and click OK.</p> <p>To delete a prefix list, select it and click Delete.</p>
BGP Community	Specifies a BGP community.	<p>To add a BGP community:</p> <ol style="list-style-type: none"> 1. Click Add. 2. Enter a name for the community. 3. To add a community, click Add. 4. Enter the community ID and click OK. 5. Click OK. <p>To edit a BGP community, click Edit. Edit the settings and click OK.</p> <p>To delete a BGP community, select it and click Delete.</p>
AS Path	Specifies an AS path. This is applicable to BGP only.	<p>To add an AS path:</p> <ol style="list-style-type: none"> 1. Click Add. 2. Enter the AS path name. 3. Enter the regular expression and click OK. 4. Click OK. <p>To edit an AS path, click Edit. Edit the settings and click OK.</p> <p>To delete an AS path, select it and click Delete.</p>

Table 11: Terms Configuration Parameters

Field	Function	Your Action
Term Name	Specifies a term name.	Type or select and edit the name.

Table 11: Terms Configuration Parameters (*continued*)

Field	Function	Your Action
Source tab		
Family	Specifies an address family protocol.	Select a value from the list.
Routing Instance	Specifies a routing instance.	Select a value from the list.
RIB	Specifies the name of a routing table.	Select a value from the list
Preference	Specifies the individual preference value for the route.	Type or select and edit the value.
Metric	Specifies a metric value. You can specify up to four metric values.	Type or select and edit the value.
Interface	Specifies a 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 internal BGP (IBGP).	<p>To add an interface, select Add > Interface. Select the interface from the list. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.</p> <p>To add an address, select Add > Address. Select the address from the list.</p> <p>To remove an interface, select it and click Remove.</p>
Prefix List	Specifies a named list of IP addresses. You can specify an exact match with incoming routes.	<p>Click Add. Select the prefix list from the list and click OK.</p> <p>To remove a prefix list, select it and click Remove.</p>
Protocol	Specifies the name of the protocol from which the route was learned or to which the route is being advertised.	<p>Click Add and select the protocol from the list.</p> <p>To remove a protocol, select it and click Remove.</p>
Policy	Specifies the name of a policy to evaluate as a subroutine.	<p>Click Add. Select the policy from the list.</p> <p>To remove a policy, select it and click Remove.</p>
More	Specifies advanced configuration options for policies.	Click More for advanced configuration.
OSPF Area ID	Specifies the area identifier.	Type the IP address.
BGP Origin	Specifies the origin of the AS path information.	Select a value from the list.
Local Preference	Specifies the BGP local preference.	Type a value.
Route	Specifies the type of route.	<p>Select External.</p> <p>Select the OSPF type from the list.</p>

Table 11: Terms Configuration Parameters (*continued*)

Field	Function	Your Action
AS Path	Specifies the name of an AS path regular expression.	Click Add . Select the AS path from the list.
Community	Specifies the name of one or more communities.	Click Add . Select the community from the list.
Destination tab		
Family	Specifies an address family protocol.	Select a value from the list.
Routing Instance	Specifies a routing instance.	Select a value from the list.
RIB	Specifies the name of a routing table.	Select a value from the list.
Preference	Specifies the individual preference value for the route.	Type a value.
Metric	Specifies a metric value.	Type a value.
Interface	Specifies a 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 internal BGP (IBGP).	<p>To add an interface, select Add > Interface. Select the interface from the list. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.</p> <p>To add an address, select Add > Address. Select the address from the list.</p> <p>To delete an interface, select it and click Remove.</p>
Protocol	Specifies the name of the protocol from which the route was learned or to which the route is being advertised.	<p>Click Add and select the protocol from the list.</p> <p>To delete a protocol, select it and click Remove.</p>
Action tab		
Action	Specifies the action to take if the conditions match.	Select a value from the list.
Default Action	Specifies that any action that is intrinsic to the protocol is overridden. This action is also nonterminating, so that various policy terms can be evaluated before the policy is terminated.	Select a value from the list.
Next	Specifies the default control action if a match occurs, and there are no further terms in the current routing policy.	Select a value from the list.
Priority	Specifies 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.	Select a value from the list.
BGP Origin	Specifies the BGP origin attribute.	Select a value from the list.

Table 11: Terms Configuration Parameters (*continued*)

Field	Function	Your Action
AS Path Prepend	Affixes an AS number at the beginning of the AS path. The AS numbers are added after the local AS number has been added to the path. This action adds an AS number to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the affixed AS number is placed within a confederation sequence. Otherwise, the affixed AS number is placed with a nonconfederation sequence.	Enter a value.
AS Path Expand	Extracts the last AS number in the existing AS path and affixes that AS number to the beginning of the AS path n times, where n 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.	Select the type and type a value.
Load Balance Per Packet	Specifies that all next-hop addresses in the forwarding table must be installed and have the forwarding table perform per-packet load balancing. This policy action allows you to optimize VPLS traffic flows across multiple paths.	Select the check box to enable the option.
Tag	Specifies the tag value. The tag action sets the 32-bit tag field in OSPF external link-state advertisement (LSA) packets.	Select the action and type a value.
Metric	Changes the metric (MED) value by the specified negative or positive offset. This action is useful only in an external BGP (EBGP) export policy.	Select the action and type a value.
Route	Specifies whether the route is external.	Select the External check box to enable the option, and select the OSPF type.
Preference	Specifies the preference value.	Select the preference action and type a value.
Local Preference	Specifies the BGP local preference attribute.	Select the action and type a value.
Class of Service	<p>Specifies and applies the class-of-service parameters to routes installed into the routing table.</p> <ul style="list-style-type: none"> Source class The value entered here maintains the packet counts for a route passing through your network, based on the source address. Destination class The value entered here maintains packet counts for a route passing through your network, based on the destination address in the packet. Forwarding class 	<p>Type the source class.</p> <p>Type the destination class.</p> <p>Type the forwarding class.</p>

- Related Documentation**
- [Configuring BGP Sessions \(J-Web Procedure\) on page 11](#)
 - [Configuring an OSPF Network \(J-Web Procedure\) on page 15](#)
 - [Configuring a RIP Network \(J-Web Procedure\) on page 19](#)
 - [Configuring Static Routing \(J-Web Procedure\) on page 24](#)
 - [Layer 3 Protocols Supported on EX Series Switches on page 3](#)

Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks so that other processes on the EX Series switch can more optimally direct their resources.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and either the access interfaces (on EX3200, EX4200, and EX4500 switches) or the line cards (on EX6200 and EX8200 switches) for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for LACP packets only.



BEST PRACTICE: We recommend that, generally, you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

This topic describes:

- [Disabling or Enabling Distributed Periodic Packet Management Globally on page 31](#)
- [Disabling or Enabling Distributed Periodic Packet Management for LACP Packets on page 32](#)

Disabling or Enabling Distributed Periodic Packet Management Globally

Distributed PPM is enabled by default. Disable distributed PPM if you need to move all PPM processing to the Routing Engine. Enable distributed PPM if it was previously disabled and you need to run distributed PPM.

To disable distributed PPM:

```
[edit routing-options]
user@switch# set ppm no-delegate-processing
```

To enable distributed PPM if it was previously disabled:

```
[edit routing-options]
```

```
user@switch# delete ppm no-delegate-processing
```

Disabling or Enabling Distributed Periodic Packet Management for LACP Packets

Distributed PPM is enabled by default. Disable distributed PPM for only LACP packets if you need to move all PPM processing for LACP packets to the Routing Engine.

To disable distributed PPM for LACP packets:

```
[edit protocols]
user@switch# set lacp ppm centralized
```

To enable distributed PPM for LACP packets if it was previously disabled:

```
[edit protocols]
user@switch# delete lacp ppm centralized
```

Related Documentation

- [Understanding Distributed Periodic Packet Management on EX Series Switches on page 6](#)
- [Understanding Aggregated Ethernet Interfaces and LACP](#)

Using IPsec to Secure OSPFv3 Networks (CLI Procedure)

OSPF version 3 (OSPFv3) does not have a built-in authentication method and relies on IP Security (IPsec) to provide this functionality. You can use IPsec to secure OSPFv3 interfaces on EX Series switches.

This topic includes:

- [Configuring Security Associations on page 32](#)
- [Securing OPSFv3 Networks on page 33](#)

Configuring Security Associations

When you configure a security association (SA), include your choices for authentication, encryption, direction, mode, protocol, and security parameter index (SPI).

To configure a security association:

1. Specify a name for the security association:

```
[edit security ipsec]
user@switch# set security-association sa-name
```

2. Specify the mode of the security association:

```
[edit security ipsec security-association sa-name]
user@switch# set mode transport
```

3. Specify the type of security association:

```
[edit security ipsec security-association sa-name]
user@switch# set type manual
```

4. Specify the direction of the security association:

```
[edit security ipsec security-association sa-name]
user@switch# set direction bidirectional
```

5. Specify the value of the security parameter index:

```
[edit security ipsec security-association sa-name]  
user@switch# set spi spi-value
```

6. Specify the type of authentication to be used:

```
[edit security ipsec security-association sa-name]  
user@switch# set authentication algorithm type
```

7. Specify the encryption algorithm and key:

```
[edit security ipsec security-association sa-name]  
user@switch# set encryption algorithm algorithm key type
```

Securing OPSFv3 Networks

You can secure the OSPFv3 network by applying the SA to the OSPFv3 configuration.

To secure the OSPFv3 network:

```
[edit protocols ospf3 area area-number interface interface-name]  
user@switch# set ipsec-sa sa-name
```

Related Documentation

- [Understanding IPsec Authentication for OSPF Packets on EX Series Switches on page 7](#)
- [Configuring an OSPF Network \(J-Web Procedure\) on page 15](#)
- [Junos OS System Basics Configuration Guide](#)

CHAPTER 3

Verifying Layer 3 Protocols Configuration

- [Monitoring BGP Routing Information on page 35](#)
- [Monitoring OSPF Routing Information on page 37](#)
- [Monitoring RIP Routing Information on page 40](#)
- [Monitoring Routing Information on page 41](#)

Monitoring BGP Routing Information

Purpose	Use the monitoring functionality to monitor BGP routing information on the routing device.
Action	<p>To view BGP routing information in the J-Web interface, select Monitor > Routing > BGP Information.</p> <p>To view BGP routing information in the CLI, enter the following commands:</p> <ul style="list-style-type: none">• show bgp summary• show bgp neighbor
Meaning	Table 12 on page 35 summarizes key output fields in the BGP routing display in the J-Web interface.

Table 12: Summary of Key BGP Routing Output Fields

Field	Values	Additional Information
BGP Peer Summary		
Total Groups	Number of BGP groups.	
Total Peers	Number of BGP peers.	
Down Peers	Number of unavailable BGP peers.	
Unconfigured Peers	Address of each BGP peer.	
RIB Summary tab		
RIB Name	Name of the RIB group.	

Table 12: Summary of Key BGP Routing Output Fields (*continued*)

Field	Values	Additional Information
Total Prefixes	Total number of prefixes from the peer, both active and inactive, that are in the routing table.	
Active Prefixes	Number of prefixes received from the EBGp peers that are active in the routing table.	
Suppressed Prefixes	Number of routes received from EBGp peers currently inactive because of damping or other reasons.	
History Prefixes	History of the routes received or suppressed.	
Dumped Prefixes	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	
Pending Prefixes	Number of pending routes.	
State	Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete.	
BGP Neighbors		
Details	Click this button to view the selected BGP neighbor details.	
Peer Address	Address of the BGP neighbor.	
Autonomous System	AS number of the peer.	

Table 12: Summary of Key BGP Routing Output Fields (*continued*)

Field	Values	Additional Information
Peer State	<p>Current state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a TCP connection in an attempt to connect to a peer. If the connection is successful, BGP sends an open message. • Connect—BGP is waiting for the TCP connection to become complete. • Established—The BGP session has been established, and the peers are exchanging BGP update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer. 	<p>Generally, the most common states are Active, which indicates a problem establishing the BGP connection, and Established, which indicates a successful session setup. The other states are transition states, and BGP sessions normally do not stay in those states for extended periods of time.</p>
Elapsed Time	Elapsed time since the peering session was last reset.	
Description	Description of the BGP session.	

- Related Documentation**
- [Configuring BGP Sessions \(J-Web Procedure\) on page 11](#)
 - [Layer 3 Protocols Supported on EX Series Switches on page 3](#)

Monitoring OSPF Routing Information

- Purpose** Use the monitoring functionality to monitor OSPF routing information on routing devices.
- Action** To view OSPF routing information in the J-Web interface, select **Monitor > Routing > OSPF Information**.
- To view OSPF routing information in the CLI, enter the following CLI commands:
- **show ospf neighbor**
 - **show ospf interface**
 - **show ospf statistics**
- Meaning** [Table 13 on page 38](#) summarizes key output fields in the OSPF routing display in the J-Web interface.

Table 13: Summary of Key OSPF Routing Output Fields

Field	Values	Additional Information
OSPF Interfaces		
Interface	Name of the interface running OSPF.	
State	State of the interface: BDR , Down , DR , DRother , Loop , PtToPt , or Waiting .	The Down state, indicating that the interface is not functioning, and PtToPt state, indicating that a point-to-point connection has been established, are the most common states.
Area	Number of the area that the interface is in.	
DR ID	Address of the area's designated device.	
BDR ID	Address of the area's backup designated device.	
Neighbors	Number of neighbors on this interface.	
Adjacency Count	Number of devices in the area using the same area identifier.	
Stub Type	The areas into which OSPF does not flood AS external advertisements	
Passive Mode	In this mode the interface is present on the network but does not transmit or receive packets.	
Authentication Type	The authentication scheme for the backbone or area.	
Interface Address	The IP address of the interface.	
Address Mask	The subnet mask or address prefix.	
MTU	The maximum transmission unit size.	
Interface Cost	The path cost used to calculate the root path cost from any given LAN segment is determined by the total cost of each link in the path.	
Hello Interval	How often the routing device sends hello packets out of the interface.	
Dead Interval	The interval during which the routing device receives no hello packets from the neighbor.	
Retransmit Interval	The interval for which the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements to an interface's neighbors.	

Table 13: Summary of Key OSPF Routing Output Fields (*continued*)

Field	Values	Additional Information
OSPF Statistics		
Packets tab		
Sent	Displays the total number of packets sent.	
Received	Displays the total number of packets received.	
Details tab		
Flood Queue Depth	Number of entries in the extended queue.	
Total Retransmits	Number of retransmission entries enqueued.	
Total Database Summaries	Total number of database description packets.	
OSPF Neighbors		
Address	Address of the neighbor.	
Interface	Interface through which the neighbor is reachable.	
State	State of the neighbor: Attempt, Down, Exchange, ExStart, Full, Init, Loading, or 2way.	Generally, only the Down state, indicating a failed OSPF adjacency, and the Full state, indicating a functional adjacency, are maintained for more than a few seconds. The other states are transitional states that a neighbor is in only briefly while an OSPF adjacency is being established.
ID	ID of the neighbor.	
Priority	Priority of the neighbor to become the designated router.	
Activity Time	The activity time.	
Area	Area that the neighbor is in.	
Options	Option bits received in the hello packets from the neighbor.	
DR Address	Address of the designated router.	
BDR Address	Address of the backup designated router.	
Uptime	Length of time since the neighbor came up.	

Table 13: Summary of Key OSPF Routing Output Fields (*continued*)

Field	Values	Additional Information
Adjacency	Length of time since the adjacency with the neighbor was established.	

- Related Documentation**
- [Configuring an OSPF Network \(J-Web Procedure\) on page 15](#)
 - [Layer 3 Protocols Supported on EX Series Switches on page 3](#)

Monitoring RIP Routing Information

- Purpose** Use the monitoring functionality to monitor RIP routing on routing devices.
- Action** To view RIP routing information in the J-Web interface, select **Monitor > Routing > RIP Information**.
- To view RIP routing information in the CLI, enter the following CLI commands:
- **show rip statistics**
 - **show rip neighbor**
- Meaning** [Table 14 on page 40](#) summarizes key output fields in the RIP routing display in the J-Web interface.

Table 14: Summary of Key RIP Routing Output Fields

Field	Values	Additional Information
RIP Statistics		
Protocol Name	The RIP protocol name.	
Port number	The port on which RIP is enabled.	
Hold down time	The interval during which routes are neither advertised nor updated.	
Global routes learned	Number of RIP routes learned on the logical interface.	
Global routes held down	Number of RIP routes that are not advertised or updated during the hold-down interval.	
Global request dropped	Number of requests dropped.	
Global responses dropped	Number of responses dropped.	

Table 14: Summary of Key RIP Routing Output Fields (*continued*)

Field	Values	Additional Information
RIP Neighbors		
Neighbor	Name of the RIP neighbor.	This value is the name of the interface on which RIP is enabled. Click the name to see the details for this neighbor.
State	State of the RIP connection: Up or Dn (Down).	
Source Address	Local source address.	This value is the configured address of the interface on which RIP is enabled.
Destination Address	Destination address.	This value is the configured address of the immediate RIP adjacency.
Send Mode	The mode of sending RIP messages.	
Receive Mode	The mode in which messages are received.	
In Metric	Value of the incoming metric configured for the RIP neighbor.	

- Related Documentation**
- [Configuring a RIP Network \(J-Web Procedure\) on page 19](#)
 - [Layer 3 Protocols Supported on EX Series Switches on page 3](#)

Monitoring Routing Information

- Purpose** Use the monitoring functionality to view the **inet.0** routing table on the routing device.
- Action** To view the routing tables in the J-Web interface, select **Monitor > Routing > Route Information**. Apply a filter or a combination of filters to view messages. You can use filters to display relevant events.
- To view the routing table in the CLI, enter the following commands in the CLI interface:
- **show route terse**
 - **show route detail**
- Meaning** [Table 15 on page 42](#) describes the different filters, their functions, and the associated actions.
- [Table 16 on page 42](#) summarizes key output fields in the routing information display.

Table 15: Filtering Route Messages

Field	Function	Your Action
Destination Address	Specifies the destination address of the route.	Enter the destination address.
Protocol	Specifies the protocol from which the route was learned.	Enter the protocol name.
Next hop address	Specifies the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.	Enter the next hop address.
Receive protocol	Specifies the dynamic routing protocol using which the routing information was received through a particular neighbor.	Enter the routing protocol.
Best route	Specifies only the best route available.	Select the view details of the best route.
Inactive routes	Specifies the inactive routes.	Select the view details of inactive routes.
Exact route	Specifies the exact route.	Select the view details of the exact route.
Hidden routes	Specifies the hidden routes.	Select the view details of hidden routes.
Search	Applies the specified filter and displays the matching messages.	To apply the filter and display messages, click Search .

Table 16: Summary of Key Routing Information Output Fields

Field	Values	Additional Information
Static Route Addresses	The list of static route addresses.	
Protocol	Protocol from which the route was learned: Static , Direct , Local , or the name of a particular protocol.	
Preference	The preference is the individual preference value for the route.	The route preference is used as one of the route selection criteria.

Table 16: Summary of Key Routing Information Output Fields (*continued*)

Field	Values	Additional Information
Next-Hop	Network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.	<p>If a next hop is listed as Discard, all traffic with that destination address is discarded rather than routed. This value generally means that the route is a static route for which the discard attribute has been set.</p> <p>If a next hop is listed as Reject, all traffic with that destination address is rejected. This value generally means that the address is unreachable. For example, if the address is a configured interface address and the interface is unavailable, traffic bound for that address is rejected.</p> <p>If a next hop is listed as Local, the destination is an address on the host (either the loopback address or Ethernet management port 0 address, for example).</p>
Age	How long the route has been active.	
State	Flags for this route.	There are many possible flags.
AS Path	<p>AS path through which the route was learned. The letters of the AS path indicate the path origin:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete. Typically, the AS path was aggregated. 	

Related Documentation

- [Configuring Static Routing \(J-Web Procedure\) on page 24](#)
- [Configuring Static Routing \(CLI Procedure\) on page 24](#)
- [Layer 3 Protocols Supported on EX Series Switches on page 3](#)

CHAPTER 4

Configuration Statements for Layer 3 Protocols

- [\[edit protocols bgp\] Configuration Statement Hierarchy on EX Series Switches on page 45](#)
- [\[edit protocols isis\] Configuration Statement Hierarchy on EX Series Switches on page 55](#)
- [\[edit protocols ospf\] Configuration Statement Hierarchy on EX Series Switches on page 58](#)
- [\[edit protocols ospf3\] Configuration Statement Hierarchy on EX Series Switches on page 61](#)
- [\[edit protocols rip\] Configuration Statement Hierarchy on EX Series Switches on page 65](#)
- [\[edit protocols ripng\] Configuration Statement Hierarchy on EX Series Switches on page 67](#)

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

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

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

This topic lists:

- [Supported Statements in the \[edit protocols bgp\] Hierarchy Level on page 46](#)
- [Unsupported Statements in the \[edit protocols bgp\] Hierarchy Level on page 54](#)

Supported Statements in the [edit protocols bgp] Hierarchy Level

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

```
protocols {
  bgp {
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    authentication-algorithm (hmac-sha-1-96 | md5);
    authentication-key key;
    bfd-liveness-detection {
      authentication {
        algorithm algorithm-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      hold-down-interval milliseconds;
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      session-mode (automatic | multihop | single-hop);
      transmit-interval {
        threshold milliseconds;
        minimum-interval milliseconds;
      }
      version (1 | automatic);
    }
    cluster cluster-identifier;
    damping;
    description text-description;
    disable;
    export [ policy-names ];
    family inet {
      any {
        loops number;
        prefix-limit {
          maximum number;
          teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        rib-group group-name;
      }
      flow {
        loops number;
        prefix-limit {
          maximum number;
          teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        rib-group group-name;
      }
    }
  }
}
```

```

multicast {
  loops number;
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  rib-group group-name;
}
unicast {
  loops number;
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  rib-group group-name;
  topology name {
    community target identifier;
  }
}
}
family inet6 {
  any {
    loops number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
  }
  labeled-unicast {
    aggregate-label {
      community community-name;
    }
    explicit-null connected-only;
    loops number;
    per-group-label;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
    traffic-statistics {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      interval seconds;
    }
  }
}
multicast {
  loops number;
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  rib-group group-name;
}

```

```
unicast {
  loops number;
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  rib-group group-name;
  topology name {
    community target identifier;
  }
}
} # end of [edit protocols bgp family]
graceful-restart {
  disable;
  restart-time seconds;
  stale-routes-time seconds;
}
group group-name {
  advertise-external <conditional>;
  advertise-inactive;
  (advertise-peer-as | no-advertise-peer-as);
  allow [ all ip-prefix</prefix-length> ];
  as-override;
  authentication-algorithm (hmac-sha-1-96 | md5);
  authentication-key key;
  bfd-liveness-detection {
    authentication {
      algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
        meticulous-keyed-sha-1 | simple-password);
      loose-check;
    }
    holddown-interval milliseconds;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    session-mode (automatic | multihop | single-hop);
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (1 | automatic);
  }
  cluster cluster-identifier;
  damping;
  description text-description;
  export [ policy-names ];
  family inet {
    any {
      loops number;
      prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      rib-group group-name;
    }
  }
}
```

```

flow {
    loops number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
}
multicast {
    loops number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
}
unicast {
    loops number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
    topology name {
        community target identifier;
    }
}
}
family inet6 {
    any {
        loops number;
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        rib-group group-name;
    }
    labeled-unicast {
        aggregate-label {
            community community-name;
        }
        explicit-null connected-only;
        loops number;
        per-group-label;
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        rib-group group-name;
        traffic-statistics {
            file filename <files number> <size maximum-file-size> <world-readable |
                no-world-readable>;
            interval seconds;
        }
    }
}
multicast {

```

```
loops number;
prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
}
rib-group group-name;
}
unicast {
    loops number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
    topology name {
        community target identifier;
    }
}
} # end of [edit protocols bgp group family]
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
hold-time seconds;
idle-after-switch-over (seconds | forever);
import [ policy-names ];
include-mp-next-hop;
keep (all | none);
local-interface interface-name;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl tll-value;
}
neighboraddress {
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    as-override;
    authentication-algorithm algorithm;
    authentication-key key;
    bfd-liveness-detection {
        authentication {
            algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                meticulous-keyed-sha-1 | simple-password);
            loose-check;
        }
        holddown-interval milliseconds;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
```



```

session-mode (automatic | multihop | single-hop);
transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
}
version (1 | automatic);
}
cluster cluster-identifier;
damping;
description text-description;
export [ policy-names ];
family {
    (inet | inet6 | inet-mvpn | inet6-mpvn | inet-vpn | inet6-vpn | iso-vpn | l2-vpn)
    {
        (any | flow | multicast | unicast | signaling) {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            damping;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            rib-group group-name;
            topology name {
                community {
                    target identifier;
                }
            }
        }
    }
    flow {
        no-validate policy-name;
    }
    labeled-unicast {
        accepted-prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        aggregate-label {
            community community-name;
        }
        explicit-null {
            connected-only;
        }
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        resolve-vpn;
        rib inet.3;
        rib-group group-name;
        topology name {
            community {
                target identifier;
            }
        }
    }
}

```

```

    }
  }
}
route-target {
  advertise-default;
  external-paths number;
  accepted-prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
}
signaling {
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
}
}
graceful-restart {
  disable;
  restart-time seconds;
  stale-routes-time seconds;
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-interface interface-name;
local-preference preference;
log-updown;
metric-out (metric | minimum-igp <offset> | igp <offset>);
mtu-discovery;
multihop <ttl-value>;
multipath {
  multiple-as;
}
no-agggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
tcp-mss segment-size;
traceoptions {
  file filename <files number> <size size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
}
no-agggregator-id;

```

```

no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
type (external | internal);
}
hold-time seconds;
idle-after-switch-over (seconds | forever);
import [ policy-names ];
include-mp-next-hop;
keep (all | none);
local-address address;
local-as autonomous-system <loops number> < alias> <no-prepend-global-as>
    <private>;
local-interface interface-name;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl tvl-value;
}
multipath;
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
}
passive;
path-selection {
    always-compare-med;

```

```

        as-path-ignore;
    }
    peer-as autonomous-system;
    precision-timers | no-precision-timers:
    preference preference;
    remove-private;
    tcp-mss segment-size;
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}

```

Unsupported Statements in the [edit protocols bgp] Hierarchy Level

All statements in the [edit protocols bgp] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 17: Unsupported [edit protocols bgp] Configuration Statements on EX Series Switches

Statement	Hierarchy Level
NOTE: Variables, such as <i>family-name</i> , are not shown in the statements or hierarchies.	
accepted-prefix-limit	[edit protocols bgp family]
add-path	[edit protocols bgp family]
authentication-key-chain	[edit protocols bgp] [edit protocols bgp group] [edit protocols bgp group neighbor]
cisco-non-deterministic	[edit protocols bgp path-selection]
eternal-router-id	[edit protocols bgp path-selection]
igp-multiplier	[edit protocols bgp path-selection]
ipsec-sa	[edit protocols bgp] [edit protocols bgp group]
key-chain	[edit protocols bgp bfd-liveness-detection authentication] [edit protocols bgp group bfd-liveness-detection authentication] [edit protocols bgp group neighbor bfd-liveness-detection authentication]
maximum	[edit protocols bgp family accepted-prefix-limit]
med-multiplier	[edit protocols bgp path-selection]
med-plus-igp	[edit protocols bgp path-selection]

Table 17: Unsupported [edit protocols bgp] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy Level
no-validate	[edit protocols bgp family inet flow]
path-count	[edit protocols bgp family add-path send]
prefix-policy	[edit protocols bgp family add-path send]
receive	[edit protocols bgp family add-path]
send	[edit protocols bgp family add-path]
teardown	[edit protocols bgp family accepted-prefix-limit]
vpn-apply-export	[edit protocols bgp] [edit protocols bgp group neighbor]

- Related Documentation**
- BGP Configuration Guide
 - [edit protocols] Configuration Statement Hierarchy on EX Series Switches

[edit protocols isis] Configuration Statement Hierarchy on EX Series Switches

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

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

This topic lists:

- [Supported Statements in the \[edit protocols isis\] Hierarchy Level on page 55](#)
- [Unsupported Statements in the \[edit protocols isis\] Hierarchy Level on page 57](#)

Supported Statements in the [edit protocols isis] Hierarchy Level

The following hierarchy shows the [edit protocols isis] configuration statements supported on EX Series switches.

```
protocols {
  isis {
    disable;
```

```
export [ policy-names ];
graceful-restart {
  disable;
  helper-disable;
  restart-duration seconds;
}
ignore-attached-bit;
interface interface-name {
  bfd-liveness-detection {
    authentication {
      algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
        meticulous-keyed-sha-1 | simple-password);
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (0 | 1 | automatic);
  }
  checksum;
  csnp-interval (seconds | disable);
  disable;
  hello-padding (adaptive | loose | strict);
  disable;
}
level (1 | 2) {
  disable;
  hello-authentication-key key;
  hello-authentication-key-chain;
  hello-authentication-type authentication;
  hello-interval seconds;
  hold-time seconds;
  ipv4-multicast-metric number;
  ipv6-multicast-metric number;
  ipv6-unicast-metric number;
  metric metric;
  passive;
  priority number;
}
link-protection;
mesh-group (value | blocked);
no-adjacency-down-notification;
no-eligible-backup;
no-ipv4-multicast;
no-ipv6-multicast;
no-ipv6-unicast;
no-unicast-topology;
node-link-protection;
```

```

    passive;
    point-to-point;
}
level (1 | 2) {
    authentication-key key;
    authentication-type authentication;
    disable;
    external-preference preference;
    no-csnp-authentication;
    no-hello-authentication;
    no-psnp-authentication;
    preference preference;
    prefix-export-limit number;
    wide-metrics-only;
}
loose-authentication-check;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
no-ipv6-routing;
overload {
    advertise-high-metrics;
    allow-route-leaking;
    timeout seconds;
}
reference-bandwidth reference-bandwidth;
rib-group {
    inet group-name;
    inet6 group-name;
}
spf-options {
    delay milliseconds;
    holddown milliseconds;
    rapid-runs number;
}
topologies {
    ipv4-multicast;
    ipv6-multicast;
    ipv6-unicast;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
}

```

Unsupported Statements in the [edit protocols isis] Hierarchy Level

All statements in the [edit protocols isis] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 18: Unsupported [edit protocols isis] Configuration Statements on EX Series Switches

Statement	Hierarchy
NOTE: Variables, such as <i>filename</i> , are not shown in the statements or hierarchies.	
authentication-key-chain	[edit protocols isis level<1 2>]
lsp-interval	[edit protocols isis interface]
lsp-lifetime	[edit protocols isis]
key-chain	[edit protocols isis interface bfd-liveness-detection authentication]

- Related Documentation**
- IS-IS Configuration Guide
 - [edit protocols] Configuration Statement Hierarchy on EX Series Switches

[edit protocols ospf] Configuration Statement Hierarchy on EX Series Switches

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

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

This topic lists:

- [Supported Statements in the \[edit protocols ospf\] Hierarchy Level on page 58](#)
- [Unsupported Statements in the \[edit protocols ospf\] Hierarchy Level on page 61](#)

Supported Statements in the [edit protocols ospf] Hierarchy Level

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

```

protocols {
  ospf {
    area area-id {
      area-range ip-prefix</prefix-length> <exact> <override-metric metric> <restrict>;
      context-identifier
      interface interface-name {
        authentication {
          md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
          simple-password key-string;
        }
      }
    }
  }
}

```



```

}
bandwidth-based-metrics {
    bandwidth value metric number;
}
bfd-liveness-detection {
    authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
            meticulous-keyed-sha-1 | simple-password);
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    full-neighbors-only;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (0|1 | automatic);
}
dead-interval seconds;
disable;
dynamic-neighbors;
flood-reduction;
hello-interval seconds;
interface-type (nbma | p2mp | p2p);
ipsec-sa sa-name;
(link-protection | node-link-protection);
metric metric;
no-eligible-backup;
no-interface-state-traps;
no-neighbor-down-notification;
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
poll-interval seconds;
priority number;
retransmit-interval seconds;
secondary;
te-metric metric;
transit-delay seconds;
}
network-summary-export [ policy-names ];
network-summary-import [ policy-names ];
no-context-identifier-advertisement;
nssa {
    area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
    default-lsa {
        default-metric metric;
        metric-type type;
    }
}

```

```
    type-7;
  }
  (summaries | no-summaries);
}
stub <default-metric metric> <summaries | no-summaries>;
virtual-link neighbor-id router-id transit-area area-id;
}
backup-spf-options
  disable;
  downstream-paths-only;
  no-install;
}
database-protection {
  ignore-count number;
  ignore-time seconds;
  maximum-lsa number;
  reset-time seconds;
  warning-only;
  warning-threshold percent;
}
disable;
export [ policy-names ];
external-preference preference;
graceful-restart {
  disable;
  helper-disable <both | restart-signaling | standard>;
  no-strict-lsa-checking;
  notify-duration seconds;
  restart-duration seconds;
}
import [ policy-names ];
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
reference-bandwidth reference-bandwidth;
rib-group group-name;
spf-options {
  delay milliseconds;
  holddown milliseconds;
  no-ignore-our-externals;
  rapid-runs number;
}
topology {
  disable;
  rib-group group-name;
  topology-id number;
}
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
traffic-engineering {
  advertise-unnumbered-interfaces;
```

```

        credibility-protocol-preference;
        ignore-lsp-metrics;
        multicast-rpf-routes;
        no-topology;
        shortcuts <lsp-metric-into-summary>;
    }
}
}

```

Unsupported Statements in the [edit protocols ospf] Hierarchy Level

All statements in the [edit protocols ospf] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 19: Unsupported [edit protocols ospf] Configuration Statements on EX Series Switches

Statement	Hierarchy
<i>NOTE:</i> Variables, such as <i>interface-name</i> , are not shown in the statements or hierarchies.	
backup-spf-options	[edit protocols ospf topology]
key-chain	[edit protocols ospf area interface bfd-liveness-detection authentication]
overload	[edit protocols ospf topology]
prefix-export-limit	[edit protocols ospf topology]
spf-options	[edit protocols ospf topology]
topology	[edit protocols ospf area interface]

- Related Documentation**
- OSPF Configuration Guide
 - [edit protocols] Configuration Statement Hierarchy on EX Series Switches

[edit protocols ospf3] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols ospf3] hierarchy level on EX Series switches.

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

This topic lists:

- [Supported Statements in the \[edit protocols ospf3\] Hierarchy Level on page 62](#)
- [Unsupported Statements in the \[edit protocols ospf3\] Hierarchy Level on page 64](#)

Supported Statements in the [edit protocols ospf3] Hierarchy Level

The following hierarchy shows the **[edit protocols ospf3]** configuration statements supported on EX Series switches:

```
protocols {
  ospf3 {
    area area-id {
      area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
      context-identifier
      inter-area-prefix-export [ policy-names ];
      inter-area-prefix-import [ policy-names ];
      interface interface-name {
        bandwidth-based-metrics {
          bandwidth value metric number;
        }
        bfd-liveness-detection {
          authentication {
            algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
              meticulous-keyed-sha-1 | simple-password);
            loose-check;
          }
          detection-time {
            threshold milliseconds;
          }
          full-neighbors-only;
          minimum-interval milliseconds;
          minimum-receive-interval milliseconds;
          multiplier number;
          no-adaptation;
          transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
          }
          version (0|1 | automatic);
        }
        dead-interval seconds;
        disable;
        flood-reduction;
        hello-interval seconds;
        interface-type (p2mp-over-lan | p2p);
        ipsec-sa sa-name;
        (link-protection | node-link-protection);
        metric metric;
        no-eligible-backup;
        passive {
          traffic-engineering {
            remote-node-id address;
          }
        }
      }
    }
  }
}
```

```

    priority number;
    retransmit-interval seconds;
    secondary;
    transit-delay seconds;
}
no-context-identifier-advertisement;
nssa {
    area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
    default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
    }
    (summaries | no-summaries);
}
stub <default-metric metric> <summaries | no-summaries>;
}
backup-spf-options (disable | downstream-paths-only | no-install);
database-protection {
    ignore-count number;
    ignore-time seconds;
    maximum-lsa number;
    reset-time seconds;
    warning-only;
    warning-threshold percent;
}
disable;
export [ policy-names ];
external-preference preference;
graceful-restart {
    disable;
    helper-disable;
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
import [ policy-names ];
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
realm (ipv4-multicast| ipv6-multicast) {
    ... same statements as at the [edit protocols ospf3] hierarchy level, EXCEPT FOR ...
    area area-id {
        interface interface-name {
            no-eligible-backup; # NOT valid at this level
        }
    }
    backup-spf-options { ... } # NOT valid at this level
    realm realm-identifier { ... } # NOT valid at this level
    traffic-engineering { ... } # NOT valid at this level
}
realm ipv4-unicast {
    ... same statements as at the [edit protocols ospf3] hierarchy level, PLUS ...
    area area-id {

```

```

        interface interface-name {
        }
    }
    realm ipv6-unicast {
    ... same statements as at the [edit protocols ospf3] hierarchy level, PLUS ...
    disable;
    backup-spf-options {
        disable;
        downstream-paths-only;
        no-install;
    }
    }
    reference-bandwidth reference-bandwidth;
    rib-group group-name;
    spf-options {
        delay milliseconds;
        holddown milliseconds;
        no-ignore-our-externals;
        rapid-runs number;
    }
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    traffic-engineering {
        ignore-lsp-metrics;
        shortcuts <lsp-metric-into-summary>;
    }
    }
}

```

Unsupported Statements in the [edit protocols ospf3] Hierarchy Level

All statements in the [edit protocols ospf3] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exception:

Table 20: Unsupported [edit protocols ospf 3] Configuration Statements on EX Series Switches

Statement	Hierarchy
NOTE: Variables, such as <i>interface-name</i> , are not shown in the statements or hierarchies.	
key-chain	[edit protocols ospf3 area interface bfd-liveness-detection authentication]

- Related Documentation**
- OSPF Configuration Guide
 - [edit protocols] Configuration Statement Hierarchy on EX Series Switches

[edit protocols rip] Configuration Statement Hierarchy on EX Series Switches

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

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

This topic lists:

- [Supported Statements in the \[edit protocols rip\] Hierarchy Level on page 65](#)
- [Unsupported Statements in the \[edit protocols rip\] Hierarchy Level on page 67](#)

Supported Statements in the [edit protocols rip] Hierarchy Level

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

```
protocols {
  rip {
    authentication-key password;
    authentication-type type;
    (check-zero | no-check-zero);
    group group-name {
      bfd-liveness-detection {
        authentication {
          algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
            meticulous-keyed-sha-1 | simple-password);
          loose-check;
        }
        detection-time {
          threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        version (1 | automatic);
      }
      export [ policy-names ];
      import [ policy-names ];
      metric-out metric;
      neighbor neighbor-name {
```

```
any-sender;
authentication-key password;
authentication-type type;
bfd-liveness-detection {
  authentication {
    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
      meticulous-keyed-sha-1 | simple-password);
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}
(check-zero | no-check-zero);
import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
route-timeout seconds;
send (broadcast | multicast | none | version-1);
update-interval seconds;
}
preference preference;
route-timeout seconds;
update-interval seconds;
}
graceful-restart {
  disable;
  restart-time seconds;
}
holddown seconds;
import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
rib-group group-name;
route-timeout seconds;
send (broadcast | multicast | none | version-1);
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
update-interval seconds;
}
}
```


Unsupported Statements in the [edit protocols rip] Hierarchy Level

All statements in the **[edit protocols rip]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 21: Unsupported [edit protocols-rip] Configuration Statements on EX Series Switches

Statement	Hierarchy
demand-circuit	[edit protocols rip group] [edit protocols rip group neighbor]
key-chain	[edit protocols rip group bfd-liveness-detection authentication] [edit protocols rip group neighbor bfd-liveness-detection authentication]
max-retrans-time	[edit protocols rip group] [edit protocols rip group neighbor]

NOTE: Variables, such as *group-name*, are not shown in the statements or hierarchies.

- Related Documentation**
- RIP Configuration Guide
 - [edit protocols] Configuration Statement Hierarchy on EX Series Switches

[edit protocols ripng] Configuration Statement Hierarchy on EX Series Switches

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

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

This topic lists:

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

Supported Statements in the [edit protocols ripng] Hierarchy Level

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

```
protocols {
  ripng {
```

```
graceful-restart {
  disable;
  restart-time seconds;
}
group group-name {
  export [ policy-names ];
  import [ policy-names ];
  metric-out metric;
  neighbor neighbor-name {
    import [ policy-names ];
    metric-in metric;
    receive <none>;
    route-timeout seconds;
    send <none>;
    update-interval seconds;
  }
  preference number;
  route-timeout seconds;
  update-interval seconds;
}
holddown seconds;
import [ policy-names ];
metric-in metric;
receive <none>;
route-timeout seconds;
send <none>;
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
update-interval seconds;
}
```

Unsupported Statements in the [edit protocols ripng] Hierarchy Level

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

Related Documentation

- RIPng Configuration Guide
- [edit protocols] Configuration Statement Hierarchy on EX Series Switches


accept-remote-nexthop

Syntax	accept-remote-nexthop;
Hierarchy Level	<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>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Specify that a single-hop EBGP peer accepts a remote next hop with which it does not share a common subnet. Configure a separate import policy on the EBGP peer to specify the remote next hop. You cannot configure multihop and accept-remote-nexthop statements for the same EPBG peer.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Single-Hop EBGP Peers to Accept Remote Next Hops • Understanding Route Advertisement • multipath on page 265

active

Syntax	(active passive);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Determine whether static, aggregate, or generated routes are removed from the routing and forwarding tables when they become inactive. Static routes are only removed from the routing table if the next hop becomes unreachable. This can occur if the local or neighbor interface goes down. Routes that have been configured to remain continually installed in the routing and forwarding tables are marked with reject next hops when they are inactive.</p> <ul style="list-style-type: none"> • active—Remove a route from the routing and forwarding tables when it becomes inactive. • passive—Have a route remain continually installed in the routing and forwarding tables even when it becomes inactive. <p>Include the active statement when configuring an individual route in the route portion of the static statement to override a passive option specified in the defaults portion of the statement.</p>
Default	active
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Examples: Configuring Static Routes • Example: Summarizing Routes Through Route Aggregation • Example: Conditionally Generating Static Routes

advertise-external

Syntax	<code>advertise-external {conditional};</code>
Hierarchy Level	<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 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 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 group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify BGP to advertise the best external route into an IBGP mesh group, a route reflector cluster, or an AS confederation even if the best route is an internal route. In order to configure the advertise-external statement on a route reflector, you must disable intracluster reflection with the no-client-reflect statement. The advertise-external statement is supported at both the group and neighbor level. If you configure the statement at the neighbor level, you must configure it for all neighbors in a group. Otherwise, the group is automatically split into different groups.</p>
	<div>  <p>NOTE: When configuring the advertise-external statement for an AS confederation, it is recommended that EBGP peers belonging to different autonomous systems are configured in a separate EBGP peer group. This ensures consistency while BGP sends the best external route to peers in the configured peer group.</p> </div>
Options	<p>conditional—(Optional) Advertise the best external path only if the route selection process reaches the point at which the multiple exit discriminator (MED) metric is evaluated. The conditional option restricts advertisement to when the best external path and the active path are equal until the MED step of the route selection process. This implies that external routes with a longer AS path length than the active path, for instance, are not advertised. The criteria used for selecting the best external path is the same whether or not the conditional option is configured.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BGP Route Advertisement • Understanding Route Advertisement

- [advertise-inactive on page 72](#)

advertise-inactive

Syntax	advertise-inactive;
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [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>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</pre>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	BGP will be the best advertised route even if the routing table does not select it as an active route.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring the Preference Value for BGP Routes• Example: Configuring BGP Route Preference (Administrative Distance)• Understanding Route Advertisement• advertise-external on page 71

advertise-peer-as

Syntax	advertise-peer-as;
Hierarchy Level	<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 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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Disable the default behavior of suppressing AS routes.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BGP Route Advertisement • Understanding Route Advertisement • no-advertise-peer-as

aggregate (Routing)

Syntax	<pre>aggregate { defaults { ... aggregate-options ... } route destination-prefix { policy policy-name; ... aggregate-options ... } }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit routing-options],</p> <p>[edit routing-options rib <i>routing-table-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Configure aggregate routes.
Options	<p>aggregate-options—Additional information about aggregate routes that is included with the route when it is installed in the routing table. Specify zero or more of the following options in aggregate-options. Each option is explained separately.</p> <ul style="list-style-type: none">• (active passive);• as-path <<i>as-path</i>> <origin (egp igp incomplete)> <atomic-aggregate> <aggregator <i>as-number</i> <i>ip-address</i>>;• (brief full);• community [<i>community-ids</i>];• discard;• (metric metric2 metric3 metric4) <i>value</i> <type <i>type</i>>;• (preference preference2 color color2) <i>preference</i> <type <i>type</i>>;• tag string; <p>defaults—Specify global aggregate route options. These options only set default attributes inherited by all newly created aggregate routes. These are treated as global defaults</p>

and apply to all the aggregate routes you configure in the **aggregate** statement. This part of the **aggregate** statement is optional.

route *destination-prefix*—Configure a nondefault aggregate route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.



The **policy** 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	<ul style="list-style-type: none"> • Example: Summarizing Routes Through Route Aggregation


aggregate-label

Syntax	<pre>aggregate-label { community <i>community-name</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp family inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet-vpn unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet-vpn6 unicast], [edit protocols bgp family inet labeled-unicast], [edit protocols bgp family inet6 labeled-unicast], [edit protocols bgp family inet-vpn unicast], [edit protocols bgp family inet6-vpn unicast]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Specify matching criteria (in the form of a community) such that all routes which match are assigned the same VPN label, selected from one of the several routes in the set defined by this criteria. This reduces the number of VPN labels that the router must consider, and aggregates the received labels.
Options	community <i>community-name</i> —Specify the name of the community to which to apply the aggregate label.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Aggregate Labels for VPNs

allow

Syntax	allow (all [<i>network/mask-length</i>]);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Implicitly configure BGP peers, allowing peer connections from any of the specified networks or hosts. To configure multiple BGP peers, configure one or more networks and hosts within a single allow statement or include multiple allow statements.
	<div> NOTE: You cannot define a BGP group with dynamic peers with BGP authentication enabled.</div>
Options	all —Allow all addresses, which is equivalent to 0.0.0.0/0 (or ::/0). <i>network/mask-length</i> —IPv6 or IPv4 network number of a single address or a range of allowable addresses for BGP peers, followed by the number of significant bits in the subnet mask.
	<div> NOTE: You cannot define a BGP group with dynamic peers with authentication enabled.</div>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• neighbor on page 266

any-sender

Syntax	<code>any-sender;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rip 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 rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Disable strict sender address checks.</p> <p>If the sender of a RIP message does not belong to the subnet of the interface, the message is discarded. This situation might cause problems with dropped packets when RIP is running on point-to-point interfaces, or when the addresses on the interfaces do not fall in the same subnet. You can resolve this by disabling strict address checks on the RIP traffic.</p>
	<div>  <p>NOTE: The <code>any-sender</code> statement is supported only for peer-to-peer interfaces.</p> </div>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring RIP

area

Syntax

```
area area-id {  
    interface interface-name {  
        passive;  
        topology (ipv4-multicast | name) {  
            disable;  
        }  
    }  
    virtual-link neighbor-id router-id transit-area area-id {  
        topology (ipv4-multicast | name) {  
            disable;  
        }  
    }  
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast |
 ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast |
 ipv4-multicast | ipv6-multicast)]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for the **realm** statement introduced in Junos OS Release 9.2.
Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Specify the area identifier for this routing device to use when participating in OSPF routing. All routing devices in an area must use the same area identifier to establish adjacencies.

Specify multiple **area** statements to configure the routing device as an area border router. An area border router does not automatically summarize routes between areas. Use the **area-range** statement to configure route summarization. By definition, an area border router must be connected to the backbone area either through a physical link or through a virtual link. To create a virtual link, include the **virtual-link** statement.

To specify that the routing device is directly connected to the OSPF backbone, include the **area 0.0.0.0** statement.

All routing devices on the backbone must be contiguous. If they are not, use the **virtual-link** statement to create the appearance of connectivity to the backbone.

You can also configure any interface that belongs to one or more topologies to advertise the direct interface addresses without actually running OSPF on that interface. By default, OSPF must be configured on an interface in order for direct interface addresses to be advertised as interior routes.



NOTE: If you configure an interface with the `passive` statement, it applies to all the topologies to which the interface belongs. You cannot configure an interface as passive for only one specific topology and have it remain active for any other topologies to which it belongs.

Options	<i>area-id</i> —Area identifier. The identifier can be up to 32 bits. It is common to specify the area number as a simple integer or an IP address. Area number 0.0.0.0 is reserved for the OSPF backbone area.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • OSPF Areas and Router Functionality Overview • Understanding Multiple Address Families for OSPFv3 • virtual-link on page 417

area-range

Syntax	area-range <i>network/mask-length</i> <exact> <override-metric <i>metric</i> > <restrict>;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm 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>
Description	<p>(Area border routers only) For an area, summarize a range of IP addresses when sending summary link advertisements (within an area). To summarize multiple ranges, include multiple area-range statements.</p> <p>For a not-so-stubby area (NSSA), summarize a range of IP addresses when sending NSSA link-state advertisements. The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas. To specify multiple prefixes, include multiple area-range statements. All external routes learned within the area that do not fall into one of the prefixes are advertised individually to other areas.</p>
Default	By default, area border routers do not summarize routes being sent from one area to other areas, but rather send all routes explicitly.
Options	<p>exact—(Optional) Summarization of a route is advertised only when an exact match is made with the configured summary range.</p> <p>mask-length—Number of significant bits in the network mask.</p> <p>network—IP address. You can specify one or more IP addresses.</p>

override-metric *metric*—(Optional) Override the metric for the IP address range and configure a specific metric value.


restrict—(Optional) Do not advertise the configured summary. This hides all routes that are contained within the summary, effectively creating a route filter.

Range: 1 through 16,777,215

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

Related Documentation	<ul style="list-style-type: none">• Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements
------------------------------	---

as-override

Syntax	as-override;
Hierarchy Level	<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 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 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 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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Compare the AS path of an incoming advertised route with the AS number of the BGP peer under the group and replace all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer.
<div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;">  </div> <div> <p>NOTE: The as-override statement is specific to a particular BGP group. This statement does not affect peers from the same remote AS configured in different groups.</p> </div> </div>	
<p>Enabling the AS override feature allows routes originating from an AS to be accepted by a router residing in the same AS. Without AS override enabled, the routing device refuses the route advertisement once the AS path shows that the route originated from its own AS. This is done by default to prevent route loops. The as-override statement overrides this default behavior.</p> <p>Note that enabling the AS override feature may result in routing loops. Use this feature only for specific applications that require this type of behavior, and in situations with strict network control. One application is the IGP protocol between the provider edge routing device and the customer edge routing device in a virtual private network.</p>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring a Layer 3 VPN with Route Reflection and AS Override • Understanding Route Advertisement • Junos OS VPNs Configuration Guide

as-path (Routing Options)

Syntax	<code>as-path <as-path> <aggregator as-number ip-address> <atomic-aggregate> <origin (egp igp incomplete)>;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)], [edit routing-options (aggregate generate static) (defaults route)], [edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Associate BGP autonomous system (AS) path information with a static, aggregate, or generated route.</p> <p>In Junos OS Release 9.1 and later, the numeric range for the AS number is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value in the range from 0.0 through 65535.65535 in AS-dot notation format.</p>
Default	No AS path information is associated with static routes.
Options	<p>aggregator—(Optional) Attach the BGP aggregator path attribute to the aggregate route. You must specify the last AS number that formed the aggregate route (encoded as two octets) for as-number, followed by the IP address of the BGP system that formed the aggregate route for ip-address.</p>

as-path—(Optional) AS path to include with the route. It can include a combination of individual AS path numbers and AS sets. Enclose sets in brackets ([]). The first AS number in the path represents the AS immediately adjacent to the local AS. Each subsequent number represents an AS that is progressively farther from the local AS, heading toward the origin of the path. You cannot specify a regular expression for **as-path**. You must use a complete, valid AS path.

atomic-aggregate—(Optional) Attach the BGP **atomic-aggregate** path attribute to the aggregate route. This path attribute indicates that the local system selected a less specific route instead of a more specific route.

origin egp—(Optional) BGP origin attribute that indicates that the path information originated in another AS.

origin igp—(Optional) BGP origin attribute that indicates that the path information originated within the local AS.

origin incomplete—(Optional) BGP origin attribute that indicates that the path information was learned by some other means.

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

Related Documentation	<ul style="list-style-type: none">• Examples: Configuring Static Routes• Example: Summarizing Routes Through Route Aggregation• Example: Conditionally Generating Static Routes• Understanding a 4-Byte Capable Router AS Path Through a 2-Byte Capable Domain in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>
------------------------------	--

asm-override-ssm

Syntax	asm-override-ssm;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Enable the routing device to accept any-source multicast join messages (*G) for group addresses that are within the default or configured range of source-specific multicast groups.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Source-Specific Multicast Groups with Any-Source Override


authentication-algorithm

Syntax	<code>authentication-algorithm <i>algorithm</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ldp session <i>session-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 <i>group</i> <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 <i>group</i> <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>ldp session <i>session-address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group</i> <i>group-name</i>],</code> <code>[edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</code> <code>[edit protocols ldp session <i>session-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 <i>group</i> <i>group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i></code> <code>neighbor <i>address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp session <i>session-address</i>]</code>
Release Information	Statement introduced in Junos OS Release 7.6. Statement introduced for BGP in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure an authentication algorithm type.
Options	<i>algorithm</i> —Specify one of the following types of authentication algorithms: <ul style="list-style-type: none">• aes-128-cmac-96—Cipher-based message authentication code (AES128, 96 bits).• hmac-sha-1-96—Hash-based message authentication code (SHA1, 96 bits).• md5—Message digest 5.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Route Authentication• Example: Configuring Route Authentication for BGP

authentication-key (Protocols BGP)

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name neighbor 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 <i>group 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 neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp group <i>group-name neighbor 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 <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name neighbor address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure an MD5 authentication key (password). Neighboring routing devices use the same password to verify the authenticity of BGP packets sent from this system.
Options	<i>key</i> —Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Route Authentication for BGP

authentication-key (Protocols IS-IS)

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>], [edit protocols isis level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Authentication key (password). Neighboring routing devices use the password to verify the authenticity of packets sent from this interface. For the key to work, you also must include the authentication-type statement.</p> <p>All routing devices must use the same password. If you are using the Junos OS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces adjacent to the Juniper Networks routing device.</p>
Default	If you do not include this statement and the authentication-type statement, IS-IS authentication is disabled.
Options	key —Authentication password. The password can be up to 1024 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
<div><div>CAUTION: A simple password for authentication is truncated if it exceeds 254 characters.</div></div>	
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring Hitless Authentication Key Rollover for IS-IS

authentication-key (Protocols RIP)

Syntax	<code>authentication-key password;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols <i>rip</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <i>neighbor neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>rip</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> <i>neighbor neighbor-name</i>],</p> <p>[edit protocols <i>rip</i>],</p> <p>[edit protocols rip group <i>group-name</i> <i>neighbor neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>rip</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor neighbor-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Require authentication for RIP route queries received on an interface.
Options	<i>password</i> —Authentication password. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Route Authentication for RIP

authentication-key-chain (Protocols BGP)

Syntax	<code>authentication-key-chain <i>key-chain</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor 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 <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name neighbor address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp <i>group group-name neighbor 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 <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for BGP, you cannot commit the <code>0.0.0.0/allow</code> statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.
Options	<i>key-chain</i> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Route Authentication for BGP• Example: Configuring BFD Authentication for Static Routes• Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols

authentication-key-chains

Syntax	<pre> authentication-key-chains { key-chain <i>key-chain-name</i> { description <i>text-string</i>; key <i>key</i> { algorithm (md5 hmac-sha-1); options (basic isis-enhanced); secret <i>secret-data</i>; start-time <i>yyyy-mm-dd.hh:mm:ss</i>; } tolerance <i>seconds</i>; } } </pre>
Hierarchy Level	[edit security]
Release Information	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure authentication key updates for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol. When the authentication-key-chains statement is configured at the [edit security] hierarchy level, and is associated with the BGP, LDP, or IS-IS protocols at the [edit protocols] hierarchy level or with the BFD protocol using the bfd-liveness-detection statement, authentication key updates can occur without interrupting routing and signaling protocols such as Open Shortest Path First (OSPF) and Resource Reservation Setup Protocol (RSVP).</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols Example: Configuring BFD Authentication for Static Routes Example: Configuring Hitless Authentication Key Rollover for IS-IS

authentication-type (Protocols IS-IS)

Syntax	<code>authentication-type <i>authentication</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable authentication and specify the authentication scheme for IS-IS. If you enable authentication, you must specify a password by including the authentication-key statement.
Default	If you do not include this statement and the authentication-key statement, IS-IS authentication is disabled.
Options	<i>authentication</i> —Authentication scheme: <ul style="list-style-type: none">• md5—Use HMAC authentication in combination with MD5. HMAC-MD5 authentication is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>.• simple—Use a simple password for authentication. The password is included in the transmitted packet, making this method of authentication relatively insecure. We recommend that you <i>not</i> use this authentication method.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Hitless Authentication Key Rollover for IS-IS• authentication-key on page 88• no-authentication-check on page 273

authentication-type (Protocols RIP)

Syntax	<code>authentication-type type;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rip],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip 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 rip],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols rip],</p> <p>[edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure the type of authentication for RIP route queries received on an interface.
Default	If you do not include this statement and the authentication-key statement, RIP authentication is disabled.
Options	<p>type—Authentication type:</p> <ul style="list-style-type: none"> • md5—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing device uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme. • none—Disable authentication. If none is configured, the configured authentication key is ignored. • simple—Use a simple password. The password is included in the transmitted packet, which makes this method of authentication relatively insecure. The password can be from 1 through 16 contiguous letters or digits long.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Route Authentication for RIP • authentication-key on page 89

autonomous-system

Syntax	<code>autonomous-system <i>autonomous-system</i> <asdot-notation> <loops <i>number</i>> { independent-domain <no-attrset>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. asdot-notation option introduced in Junos OS Release 9.3. asdot-notation option introduced in Junos OS Release 9.3 for EX Series switches. no-attrset option introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Specify the routing device's AS number.</p> <p>An autonomous system (AS) is a set of routing devices that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it. ASs are identified by a number that is assigned by the Network Information Center (NIC) in the United States (http://www.isi.edu).</p> <p>If you are using BGP on the routing device, you must configure an AS number.</p> <p>The AS path attribute is modified when a route is advertised to an EBGP peer. Each time a route is advertised to an EBGP peer, the local routing device prepends its AS number to the existing path attribute, and a value of 1 is added to the AS number.</p> <p>In Junos OS Release 9.1 and later, the numeric range is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p>
Options	<i>autonomous-system</i> —AS number. Use a number assigned to you by the NIC.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format for 4-byte AS numbers

In this example, the 4-byte AS number 65,546 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 65546;
}
```

Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte numbers

In this example, 1.10 is the AS-dot notation format for 65,546:

```
[edit]
routing-options {
  autonomous-system 1.10;
}
```

Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

In this example, the 2-byte AS number 60,000 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 60000;
}
```

asdot-notation—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

Default: Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

loops number—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

Range: 1 through 10

Default: 1



NOTE: When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the loops statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.

Use the **independent-domain** option if the loops statement must be enabled only on a subset of routing instances.

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	<ul style="list-style-type: none">• Examples: Configuring External BGP Peering• Examples: Configuring Internal BGP Peering• 4-Byte Autonomous System Numbers Overview in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>• Juniper Networks Implementation of 4-Byte Autonomous System Numbers in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>• Configuring 4-Byte Autonomous System Numbers in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>

backup-pe-group

Syntax	<pre>backup-pe-group group-name { backups [addresses]; local-address address; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.5 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.
Description	Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.
Options	<p>backups <i>addresses</i>—Specify the address of backup PE routers for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p>local-address <i>address</i>—Specify the address of the local PE router for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p>pe-group-name—Specify the name for the group of PE routers that provide ingress PE router redundancy for point-to-multipoint LSPs.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Ingress PE Redundancy• Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs

backups

Syntax	<code>backups [<i>addresses</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast backup-pe-group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast backup-pe-group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast backup-pe-group <i>group-name</i>],</p> <p>[edit routing-options multicast backup-pe-group <i>group-name</i>]</p>
Release Information	<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>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure the address of backup PEs for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.
Options	<i>addresses</i> —Addresses of other PEs in the backup group.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Ingress PE Redundancy

bandwidth (Multicast Flow Map)

Syntax	<code>bandwidth (<i>bps</i> <i>adaptive</i>);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map], [edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map], [edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map], [edit routing-options multicast flow-map]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure the bandwidth property for multicast flow maps.
Options	adaptive —Specify that the bandwidth is measured for the flows that are matched by the flow map. bps —Bandwidth, in bits per second, for the flow map. Range: 0 through any amount of bandwidth Default: 2 Mbps
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring a Multicast Flow Map

bandwidth-based-metrics

Syntax	<pre>bandwidth-based-metrics { bandwidth <i>value</i>; metric <i>number</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p>
Description	<p>Specify a set of bandwidth threshold values and associated metric values for an OSPF interface or for a topology on an OSPF interface. When the bandwidth of an interface changes, Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value.</p>
Options	<p>bandwidth <i>value</i>—Specify the bandwidth threshold in bits per second.</p> <p>Range: 9600 through 1,000,000,000,000,000</p> <p>metric <i>number</i>—Specify a metric value to associate with a specific bandwidth value.</p> <p>Range: 1 through 65,535</p>



NOTE: You must also configure a static metric value for the OSPF interface or topology with the metric statement. Junos OS uses this value to calculate the cost of a route from the OSPF interface or topology if the bandwidth for the interface is higher than of any bandwidth threshold values configured for bandwidth-based metrics.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth• metric on page 249• Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth

bfd-liveness-detection (Protocols BGP)

Syntax `bfd-liveness-detection {`
 `authentication {`
 `algorithm` *algorithm-name*;
 `key-chain` *key-chain-name*;
 `loose-check`;
 `}`
 `detection-time {`
 `threshold` *milliseconds*;
 `}`
 `hold-down-interval` *milliseconds*;
 `minimum-interval` *milliseconds*;
 `minimum-receive-interval` *milliseconds*;
 `multiplier` *number*;
 `no-adaptation`;
 `session-mode` (automatic | multihop | single-hop);
 `transmit-interval {`
 `minimum-interval` *milliseconds*;
 `threshold` *milliseconds*;
 `}`
 `version` (1 | automatic);
`}`

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

Release Information Statement introduced in Junos OS Release 8.1.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
detection-time threshold and **transmit-interval threshold** options introduced in Junos OS Release 8.2
 Support for logical routers introduced in Junos OS Release 8.3.
 Support for IBGP and multihop EBGP sessions introduced in Junos OS Release 8.3.
holddown-interval statement introduced in Junos OS Release 8.5. You can configure this statement only for EBGP peers at the [edit protocols bgp group *group-name* neighbor *address*] hierarchy level.
no-adaptation statement introduced in Junos OS Release 9.0.
 Support for BFD authentication introduced in Junos OS Release 9.6.

Support for BFD on IPv6 interfaces with BGP introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure bidirectional failure detection (BFD) timers and authentication for BGP.

For IBGP and multihop EBGP support, configure the **bfd-liveness-detection** statement at the global **[edit bgp protocols]** hierarchy level. You can also configure IBGP and multihop support for a routing instance or a logical system.

Options **authentication algorithm** *algorithm-name* (Optional)—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name* (Optional)—Associate a security key with the specified BFD session using the name of the security keychain. The keychain name must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds* (Optional)—Configure a threshold. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

holddown-interval *milliseconds* (Optional)—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. When you configure the hold-down interval for the BFD protocol for EBGp, the BFD session is unaware of the BGP session during this time. In this case, if the BGP session goes down during the configured hold-down interval, BFD already assumes it is down and does not send a state change notification. The **holddown-interval** statement is supported only for EBGp peers at the **[edit protocols bgp group group-name neighbor address]** hierarchy level. If the BFD session goes down and then comes back up during the configured hold-down interval, the timer is restarted. You must configure the hold-down interval on both EBGp peers. If you configure the hold-down interval for a multihop EBGp session, you must also configure a local IP address by including the **local-address** statement at the **[edit protocols bgp group group-name]** hierarchy level.

Range: 0 through 255,000

Default: 0

minimum-interval *milliseconds* (Required)—Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately (using the **minimum-receive-interval** and **transmit-interval** statements).

Range: 1 through 255,000

minimum-receive-interval *milliseconds* (Optional)—Configure only the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Range: 1 through 255,000

multiplier *number* (Optional)—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation (Optional)—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable to not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds* (Optional)—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds* (Optional)—Configure only the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.

Range: 1 through 255,000

version (Optional)—Configure the BFD version to detect.

Range: 1 or **automatic** (autodetect the BFD version)

Default: **automatic**

The remaining statements are explained separately.

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

Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes• Example: Configuring BFD Authentication for Static Routes• Example: Configuring BFD on Internal BGP Peer Sessions• Example: Configuring BFD Authentication for BGP• Understanding BFD for BGP
------------------------------	---

bfd-liveness-detection (Protocols IS-IS)

Syntax	<pre> bfd-liveness-detection { authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check; } detection-time { threshold <i>milliseconds</i>; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } version (1 automatic); } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. detection-time threshold and transmit-interval threshold options added in Junos OS Release 8.2. Support for logical systems introduced in Junos OS Release 8.3. no-adaptation statement introduced in Junos OS Release 9.0. authentication algorithm, authentication key-chain, and authentication loose-check options introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure bidirectional failure detection timers and authentication.
Options	<p>authentication algorithm <i>algorithm-name</i> —Configure the algorithm used to authenticate the specified BFD session: simple-password, keyed-md5, keyed-sha-1, meticulous-keyed-md5, meticulous-keyed-sha-1.</p> <p>authentication key-chain <i>key-chain-name</i> —Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the authentication-key-chains key-chain statement at the [edit security] hierarchy level.</p> <p>authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.</p>

detection-time threshold *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

minimum-interval *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.

Range: 1 through 255,000

minimum-receive-interval *milliseconds*—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

Range: 1 through 255,000

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

Range: 1 through 255

Default: 3

no-adaptation—Specify that BFD sessions not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds*—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version)

Default: automatic

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

Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for IS-IS• Example: Configuring BFD Authentication for IS-IS
------------------------------	--

bfd-liveness-detection (Protocols OSPF)

Syntax	<pre> bfd-liveness-detection { authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check; } detection-time { threshold <i>milliseconds</i>; } full-neighbors-only minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } version (1 automatic); } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>detection-time threshold and transmit-interval threshold options added in Junos OS Release 8.2.</p> <p>Support for logical systems introduced in Junos OS Release 8.3.</p> <p>no-adaptation option introduced in Junos OS Release 9.0.</p> <p>no-adaptation option introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 introduced in Junos OS Release 9.3.</p> <p>Support for OSPFv3 introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>full-neighbors-only option introduced in Junos OS Release 9.5.</p> <p>full-neighbors-only option introduced in Junos OS Release 9.5 for EX Series switches.</p>

authentication algorithm, **authentication key-chain**, and **authentication loose-check** options introduced in Junos OS Release 9.6.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure bidirectional failure detection timers and authentication for OSPF.

The remaining statements are explained separately.

Options **authentication algorithm** *algorithm-name*—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

full-neighbors-only—Establish BFD sessions only for OSPF neighbors in the full state. The default behavior is to establish BFD sessions for all OSPF neighbors.

minimum-interval *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

Range: 1 through 255,000 milliseconds

minimum-receive-interval *milliseconds*—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

Range: 1 through 255,000 milliseconds

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

Range: 1 through 255

Default: 3

no-adaptation—Specify that BFD sessions should not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds*—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established

a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

Default: **automatic**

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

Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for OSPF• Example: Configuring BFD Authentication for OSPF
------------------------------	--

bfd-liveness-detection (Protocols RIP)

Syntax	<pre> bfd-liveness-detection { authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check; } detection-time { threshold <i>milliseconds</i>; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } version (1 automatic); } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>], [edit protocols rip group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.0.</p> <p>Options detection-time threshold and transmit-interval threshold introduced in Junos OS Release 8.2.</p> <p>Support for logical systems introduced in Junos OS Release 8.3.</p> <p>Option no-adaptation introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Options authentication algorithm, authentication key-chain, and authentication loose-check introduced in Junos OS Release 9.6.</p> <p>Options authentication algorithm, authentication key-chain, and authentication loose-check introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure bidirectional failure detection timers and authentication.
Options	<p>authentication algorithm <i>algorithm-name</i>—Configure the algorithm used to authenticate the specified BFD session: simple-password, keyed-md5, keyed-sha-1, meticulous-keyed-md5, or meticulous-keyed-sha-1.</p> <p>authentication key-chain <i>key-chain-name</i>—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the authentication-key-chains key-chain statement at the [edit security] hierarchy level.</p>

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication is not configured at both ends of the BFD session.

detection-time threshold *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

minimum-interval *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

Range: 1 through 255,000 milliseconds

minimum-receive-interval *milliseconds*—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

Range: 1 through 255,000 milliseconds

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

Range: 1 through 255

Default: 3

no-adaptation—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds*—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

Default: automatic

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

- Related Documentation**
- Example: Configuring BFD for RIP
 - Example: Configuring BFD Authentication for RIP

bfd-liveness-detection (Routing Options Static Route)

Syntax

```

bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    holddown-interval milliseconds;
    local-address ip-address;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-receive-ttl number;
    multiplier number;
    neighbor address;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}

```

Hierarchy Level

```

[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options rib routing-table-name static route destination-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options rib routing-table-name static route destination-prefix qualified-next-hop
 (interface-name | address)],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options static route destination-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options static route destination-prefix qualified-next-hop (interface-name |
 address)],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
 route destination-prefix],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
 route destination-prefix qualified-next-hop (interface-name | address)],
[edit logical-systems logical-system-name routing-options static route destination-prefix],
[edit logical-systems logical-system-name routing-options static route destination-prefix
 qualified-next-hop (interface-name | address)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
 route destination-prefix],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
 route destination-prefix qualified-next-hop (interface-name | address)],
[edit routing-instances routing-instance-name routing-options static route destination-prefix],
[edit routing-instances routing-instance-name routing-options static route destination-prefix
 qualified-next-hop (interface-name | address)],
[edit routing-options rib routing-table-name static route destination-prefix],
[edit routing-options rib routing-table-name static route destination-prefix qualified-next-hop
 (interface-name | address)],
[edit routing-options static route destination-prefix],

```


[edit routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)]

Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>detection-time threshold and transmit-interval threshold options introduced in Junos OS Release 8.2.</p> <p>local-address statement introduced in Junos OS Release 8.2.</p> <p>minimum-receive-ttl statement introduced in Junos OS Release 8.2.</p> <p>Support for logical routers introduced in Junos OS Release 8.3.</p> <p>holddown-interval statement introduced in Junos OS Release 8.5.</p> <p>no-adaptation statement introduced in Junos OS Release 9.0.</p> <p>Support for IPv6 static routes introduced in Junos OS Release 9.1.</p> <p>authentication algorithm, authentication key-chain, and authentication loose-check statements introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Configure bidirectional failure detection timers and authentication criteria for static routes.</p>

Options **authentication algorithm** *algorithm-name* —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name* —Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

holddown-interval *milliseconds*—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

Range: 0 through 255,000

Default: 0

local-address *ip-address*—Enable a multihop BFD session and configure the source address for the BFD session.

minimum-interval *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.

Range: 1 through 255,000

minimum-receive-interval *milliseconds*—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

Range: 1 through 255,000

minimum-receive-ttl *number*—Configure the time to live (TTL) for the multihop BFD session.

Range: 1 through 255

Default: 255

multiplier *number*—Configure number of hello packets not received by the neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

neighbor *address*—Configure a next-hop address for the BFD session for a next hop specified as an interface name.

no-adaptation—Specify for BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295

transmit-interval minimum-interval *milliseconds*—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

Default: automatic


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

Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BFD for Static Routes • Example: Configuring BFD Authentication for Static Routes
------------------------------	---

bgp

Syntax	<code>bgp { ... }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable BGP on the routing device or for a routing instance.
Default	BGP is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• BGP Configuration Guide

bgp-orf-cisco-mode

Syntax	<code>bgp-orf-cisco-mode;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</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> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options outbound-route-filter],</p> <p>[edit protocols bgp outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit routing-options outbound-route-filter]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.
	<div>  <p>NOTE: To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.</p> </div>
Default	Disabled
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- Example: Configuring BGP Prefix-Based Outbound Route Filtering


bmp

Syntax	<pre>bmp { memory limit <i>bytes</i>; station-address (<i>ip-address</i> <i>name</i>); station-port <i>port-number</i>; statistics-timeout <i>seconds</i>; }</pre>
Hierarchy Level	[edit routing-options]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure the BGP Monitoring Protocol (BMP), which enables the routing device to collect data from the BGP Adjacency-RIB-In routing tables and periodically send that data to a monitoring station.
Options	<p>memory-limit <i>bytes</i>—(Optional) Specify a threshold at which to stop collecting BMP data if the limit is exceeded.</p> <p>Default: 10 MB</p> <p>Range: 1,048,576 through 52,428,800</p> <p>station-address (<i>ip-address</i> <i>name</i>)—Specify the IP address or a valid URL for the monitoring where BMP data should be sent.</p> <p>station-port <i>port-number</i>—Specify the port number of the monitoring station to use when sending BMP data.</p> <p>statistics-timeout <i>seconds</i>—(Optional) Specify how often to send BMP data to the monitoring station.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring the BGP Monitoring Protocol

brief

Syntax	(brief full);
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate) (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)], [edit routing-options (aggregate generate) (defaults route)], [edit routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Configure all AS numbers from all contributing paths to be included in the aggregate or generated route's path.</p> <ul style="list-style-type: none"> • brief—Include only the longest common leading sequences from the contributing AS paths. If this results in AS numbers being omitted from the aggregate route, the BGP ATOMIC_ATTRIBUTE path attribute is included with the aggregate route. • full—Include all AS numbers from all contributing paths in the aggregate or generated route's path. Include this option when configuring an individual route in the route portion of the generate statement to override a retain option specified in the defaults portion of the statement.
Default	full
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Summarizing Routes Through Route Aggregation • Example: Conditionally Generating Static Routes • aggregate on page 74 • generate on page 164

centralized

Syntax	centralized;
Hierarchy Level	[edit protocols lacp ppm]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Disable distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all PPM processing for LACP packets on the Routing Engine.</p> <p>This statement disables distributed PPM processing for only LACP packets. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the no-delegate-processing statement in the [edit routing-options ppm] hierarchy.</p> <div><p>BEST PRACTICE: We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.</p></div>
Default	Distributed PPM processing is enabled for all packets that use PPM.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure) on page 31• Configuring Aggregated Ethernet LACP (CLI Procedure)• Configuring Distributed Periodic Packet Management• Configuring Link Aggregation

check-zero

Syntax	(check-zero no-check-zero);
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols <i>rip</i>], [edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>rip</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>], [edit protocols <i>rip</i>], [edit protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>rip</i>], [edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Some of the reserved fields in RIP version 1 packets must be zero, whereas in RIP version 2 packets, most of these reserved fields can contain nonzero values. By default, RIP discards version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.</p> <p>If you find that you are receiving RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero, you can configure RIP to receive these packets even though they are being sent in violation of the specifications in RFC 1058 and RFC 2453.</p> <p>Check whether the reserved fields in a RIP packet are zero:</p> <ul style="list-style-type: none"> • check-zero—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications. • no-check-zero—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.
Default	check-zero
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring RIP

checksum (Protocols IS-IS)

Syntax	checksum;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable checksums for packets on this interface. Junos OS supports IS-IS checksums as documented in RFC 3358, <i>Optional Checksums in Intermediate System to Intermediate System (ISIS)</i> . The checksum cannot be enabled with MD5 hello authentication on the same interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Enabling Packet Checksums on IS-IS Interfaces

cluster

Syntax	<code>cluster <i>cluster-identifier</i>;</code>
Hierarchy Level	<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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.



CAUTION:

If you configure both route reflection and VPNs on the same routing device, the following modifications to the route reflection configuration cause current BGP sessions to be reset:

- Adding a cluster ID—If a BGP session shares the same AS number with the group where you add the cluster ID, all BGP sessions are reset regardless of whether the BGP sessions are contained in the same group.
- Creating a new route reflector—If you have an IBGP group with an AS number and create a new route reflector group with the same AS number, all BGP sessions in the IBGP group and the new route reflector group are reset.



NOTE: If you change the address family specified in the [edit protocols bgp family] hierarchy level, all current BGP sessions on the routing device are dropped and then reestablished.

Options *cluster-identifier*—IPv6 or IPv4 address to use as the cluster identifier.

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

Related Documentation

- Example: Configuring BGP Route Reflectors
- Understanding External BGP Peering Sessions
- [no-client-reflect on page 274](#)

community (Routing Options)

Syntax	<code>community ([<i>community-ids</i>] no-advertise no-export no-export-subconfed none);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Associate BGP community information with a static, aggregate, or generated route.
Default	No BGP community information is associated with static routes.
Options	<p><i>community-ids</i>—One or more community identifiers. The <i>community-ids</i> format varies according to the type of attribute that you use.</p> <p>The BGP community attribute format is <i>as-number:community-value</i>:</p> <ul style="list-style-type: none"> • <i>as-number</i>—AS number of the community member. It can be a value from 1 through 65,535. The AS number can be a decimal or hexadecimal value. • <i>community-value</i>—Identifier of the community member. It can be a number from 0 through 65,535. <p>For more information about BGP community attributes, see the “Configuring the Extended Communities Attribute” section in the Routing Policy Configuration Guide.</p> <p>For specifying the BGP community attribute only, you also can specify <i>community-ids</i> as one of the following well-known community names defined in RFC 1997:</p> <ul style="list-style-type: none"> • no-advertise—Routes containing this community name are not advertised to other BGP peers. • no-export—Routes containing this community name are not advertised outside a BGP confederation boundary. • no-export-subconfed—Routes containing this community name are not advertised to external BGP peers, including peers in other members’ ASs inside a BGP confederation.



NOTE: Extended community attributes are not supported at the [edit routing-options] hierarchy level. You must configure extended communities at the [edit policy-options] hierarchy level. For information about configuring extended communities, see the Routing Policy Configuration Guide.

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

- Related Documentation**
- Examples: Configuring Static Routes
 - Example: Summarizing Routes Through Route Aggregation
 - Example: Conditionally Generating Static Routes
 - [aggregate on page 74](#)
 - [generate on page 164](#)
 - [static on page 376](#)

confederation

Syntax	<code>confederation <i>confederation-autonomous-system</i> members [<i>autonomous-systems</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Specify the routing device's confederation AS number.</p> <p>If you administer multiple ASs that contain a very large number of BGP systems, you can group them into one or more <i>confederations</i>. Each confederation is identified by its own AS number, which is called a <i>confederation AS number</i>. To external ASs, a confederation appears to be a single AS. Thus, the internal topology of the ASs making up the confederation is hidden.</p> <p>The BGP path attributes NEXT_HOP, LOCAL_PREF, and MULTI_EXIT_DISC, which normally are restricted to a single AS, are allowed to be propagated throughout the ASs that are members of the same confederation.</p> <p>Because each confederation is treated as if it were a single AS, you can apply the same routing policy to all the ASs that make up the confederation.</p> <p>Grouping ASs into confederations reduces the number of BGP connections required to interconnect ASs.</p> <p>If you are using BGP, you can enable the local routing device to participate as a member of an AS confederation. To do this, include the confederation statement.</p> <p>Specify the AS confederation identifier, along with the peer AS numbers that are members of the confederation.</p> <p>Note that peer adjacencies do not form if two BGP neighbors disagree about whether an adjacency falls within a particular confederation.</p>
Options	<p><i>autonomous-systems</i>—AS numbers of the confederation members. Range: 1 through 65,535</p> <p><i>confederation-autonomous-system</i>—Confederation AS number. Use one of the numbers assigned to you by the NIC. Range: 1 through 65,535</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- Example: Configuring BGP Confederations

csnp-interval

Syntax	csnp-interval (<i>seconds</i> disable);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure the interval between complete sequence number PDUs (CSNPs) on a LAN interface.</p> <p>If the routing device is the designated router on a LAN, IS-IS sends CSN packets every 10 seconds. If the routing device is on a point-to-point interface, it sends CSN packets every 5 seconds. To protect against link-state PDU flooding, we recommend modifying the default interval.</p> <p>To modify the CSNP interval, include the csnp-interval statement.</p> <p>To configure the interface not to send any CSNPs, specify the disable option.</p>
Default	By default, IS-IS sends CSNPs periodically. If the routing device is the designated router on a LAN, IS-IS sends CSNPs every 10 seconds. If the routing device is on a point-to-point interface, it sends CSNPs every 5 seconds.
Options	<p>disable—Do not send CSNPs on this interface.</p> <p>seconds—Number of seconds between the sending of CSNPs.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 10 seconds on LAN broadcast links. 5 seconds on point-to-point links.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	• Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces

damping (Protocols BGP)

Syntax	damping;
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [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>], [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> family <i>family</i>], [edit protocols bgp], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> family <i>family</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>] [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for flap damping at the address family level introduced in Junos OS Release 12.2.</p>
Description	<p>Enable route flap damping. BGP route flapping describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability information. Flap damping reduces the number of update messages sent between BGP</p>

peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.

You typically apply flap damping to external BGP (EBGP) routes (that is, to routes in different ASs). You can also apply it within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.) The exception to this rule is when flap damping is applied at the address family level. When you apply flap damping at the address family level, it works for both IBGP and EBGP.

Default Flap damping is disabled on the routing device.

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

Related Documentation

- Examples: Configuring BGP Flap Damping
- Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family

dead-interval

Syntax	<code>dead-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</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) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Specify how long OSPF waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor.
Options	<p>seconds—Interval to wait.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: Four times the hello interval—40 seconds (broadcast and point-to-point networks); 120 seconds (nonbroadcast multiple access (NBMA) networks)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring OSPF Timers Configuring RSVP and OSPF for LMP Peer Interfaces

- [hello-interval on page 183](#)

default-lsa

Syntax	<pre>default-lsa { default-metric <i>metric</i>; metric-type <i>type</i>; type-7; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</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) area <i>area-id</i> nssa],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>On area border routers only, for a not-so-stubby area (NSSA), inject a default link-state advertisement (LSA) with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• OSPF Areas and Router Functionality Overview• Example: Configuring OSPF Not-So-Stubby Areas• nssa on page 288• stub on page 382

default-metric

Syntax	<code>default-metric <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area area-id nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area area-id stub],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> stub],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id stub],</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) area <i>area-id</i> nssa default-lsa],</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) area <i>area-id</i> stub],</p> <p>[edit protocols (ospf ospf3) area area-id nssa default-lsa],</p> <p>[edit protocols (ospf ospf3) area area-id stub],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> stub],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id stub],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> stub]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	On area border routers only, for a stub area, inject a default route with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.
Options	<p><i>metric</i>—Metric value.</p> <p>Range: 1 through 16,777,215</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Areas and Router Functionality Overview • nssa on page 288

- [stub on page 382](#)

description (Protocols BGP)

Syntax	<code>description <i>text-description</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor 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 <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name neighbor address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp <i>group group-name neighbor 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 <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Provide a description of the global, group, or neighbor configuration. If the text includes one or more spaces, enclose it in quotation marks (" "). The text is displayed in the output of the show command and has no effect on the configuration.
Options	<i>text-description</i> —Text description of the configuration. It is limited to 255 characters.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• BGP Configuration Guide

disable (Protocols BGP)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Disable BGP on the system.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• BGP Configuration Guide

disable (Protocols IS-IS)

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering],</p> <p>[edit protocols isis],</p> <p>[edit protocols isis interface <i>interface-name</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis traffic-engineering],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Disable IS-IS on the routing device, on an interface, or on a level.</p> <p>At the [edit protocols isis traffic-engineering] hierarchy level, disable IS-IS support for traffic engineering.</p> <p>Enabling IS-IS on an interface (by including the interface statement at the [edit protocols isis] or the [edit routing-instances routing-instance-name protocols isis] hierarchy level), disabling it (by including the disable statement), and not actually having IS-IS run on an interface (by including the passive statement) are mutually exclusive states.</p>
Default	<p>IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which family iso is enabled.</p> <p>IS-IS support for traffic engineering is enabled.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Multi-Level IS-IS IS-IS Overview

disable (OSPF)

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) virtual-link],</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> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</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 (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances</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 ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) virtual-link],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm 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>
Description	<p>Disable OSPF, an OSPF interface, or an OSPF virtual link.</p> <p>By default, control packets sent to the remote end of a virtual link must be forwarded using the default topology. In addition, the transit area path consists only of links that</p>

are in the default topology. You can disable a virtual link for a configured topology, but not for a default topology. Include the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id* topology name]** hierarchy level.




NOTE: If you disable the virtual link by including the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id*]** hierarchy level, you disable the virtual link for all topologies, including the default topology. You cannot disable the virtual link only in the default topology.

Default	The configured object is enabled (operational) unless explicitly disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• OSPF Configuration Overview• Configuring RSVP and OSPF for LMP Peer Interfaces

disable (Routing Options)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-options graceful-restart], [edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart], [edit routing-options graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Disable graceful restart.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Junos OS High Availability Configuration Guide

discard

Syntax	discard;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.</p> <p>To propagate static routes into the routing protocols, include the discard statement when you define the route, along with a routing policy.</p> <div style="margin-top: 20px;">  <p>NOTE: In other vendors' software, a common way to propagate static routes into routing protocols is to configure the routes so that the next-hop routing device is the loopback address (commonly, 127.0.0.1). However, configuring static routes in this way (by including a statement such as route <i>address/mask-length</i> next-hop 127.0.0.1) does not propagate the static routes, because the forwarding table ignores static routes whose next-hop routing device is the loopback address.</p> </div>
Default	When an aggregate route becomes active, it is installed in the routing table with a reject next hop, which means that ICMP unreachable messages are sent.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Summarizing Routes Through Route Aggregation • Example: Conditionally Generating Static Routes

- [aggregate on page 74](#)
- [generate on page 164](#)

domain-id

Syntax	<code>domain-id <i>domain-id</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify a domain ID for a route. The domain ID identifies the OSPF domain from which the route originated.
Options	<i>domain-id</i> —You can specify either an IP address or an IP address and a local identifier using the following format: <i>ip-address:local-identifier</i> . If you do not specify a local identifier with the IP address, the identifier is assumed to have a value of 0. Default: If the router ID is not configured in the routing instance, the router ID is derived from an interface address belonging to the routing instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Routing Between PE and CE Routers in Layer 3 VPNs

domain-vpn-tag

Syntax	<code>domain-vpn-tag <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set a virtual private network (VPN) tag for OSPFv2 external routes generated by the provider edge (PE) router.
Options	<i>number</i> —VPN tag.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring Routing Between PE and CE Routers in Layer 3 VPNs

explicit-null (Protocols BGP)


Syntax	explicit-null;
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols ldap], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols ldap], [edit protocols mpls], [edit protocols bgp <i>family</i> inet labeled-unicast], [edit protocols bgp <i>family</i> inet6 labeled-unicast], [edit protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast], [edit protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast], [edit protocols ldap], [edit routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet6 labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast], [edit routing-instances <i>instance-name</i> protocols ldap] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Advertise label 0 to the egress routing device of an LSP.

Default	If you do not include the explicit-null statement in the configuration, label 3 (implicit null) is advertised.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Advertising Explicit Null Labels to BGP Peers

export (Protocols BGP)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Apply one or more policies to routes being exported from the routing table into BGP.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Understanding Route Advertisement Routing Policy Configuration Guide import on page 193


export (Protocols IS-IS)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols isis],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>isis],</code> <code>[edit protocols isis],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols isis]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<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>NOTE: 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>
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Redistributing OSPF Routes into IS-IS• Example: Configuring an IS-IS Default Route Policy on Logical Systems

export (Protocols OSPF)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm 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>
Description	Apply one or more policies to routes being exported from the routing table into OSPF.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Understanding OSPF Routing Policy Import and Export Policies for Network Summaries Overview import on page 194 Routing Policy Configuration Guide

export (Protocols RIP)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols rip 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> rip group <i>group-name</i>],</code> <code>[edit protocols rip group <i>group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<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 metric-in and metric-out statements.</p> <div><p>NOTE: The export policy on RIP does not support manipulating routing information of the next hop.</p></div>
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring RIPimport on page 195

export (Protocols RIPng)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ripng group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i>],</p> <p>[edit protocols ripng group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>
Description	<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 export 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 metric-in and metric-out statements.</p>
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring RIPng import on page 196

export (Routing Options)

Syntax	<code>export [<i>policy-name</i>];</code>
Hierarchy Level	[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]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Apply one or more policies to routes being exported from the routing table into the forwarding table.
Options	<i>policy-name</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Load Balancing BGP Traffic• Routing Policy Configuration Guide

export-rib

Syntax	<code>export-rib <i>routing-table-name</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit routing-options rib-groups <i>group-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Specify the name of the routing table from which Junos OS should export routing information.
Options	<i>routing-table-name</i> —Routing table group name.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Exporting Specific Routes from One Routing Table Into Another Routing Table • import-rib on page 199 • passive

external-preference (Protocols IS-IS)

Syntax	external-preference <i>preference</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the preference of external routes.
Options	<i>preference</i> —Preference value. Range: 0 through 4,294,967,295 ($2^{32} - 1$) Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Route Preferences Overview• Example: Redistributing OSPF Routes into IS-IS• Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS• preference on page 314

external-preference (Protocols OSPF)

Syntax	<code>external-preference <i>preference</i>;</code>
Hierarchy Level	<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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm 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>
Description	Set the route preference for OSPF external routes.
Options	<p><i>preference</i>—Preference value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 150</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Controlling OSPF Route Preferences preference on page 315

family (Protocols BGP)

```
Syntax  family {
    (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) {
        (any | flow | labeled-unicast | multicast | unicast) {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage-threshold> idle-timeout (forever | minutes);
            }
            add-path {
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
                receive;
            }
            algp [disable];
            loops number;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            protection;
            rib-group group-name;
            topology name {
                community {
                    target identifier;
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib (inet.3 | inet6.3);
            rib-group group-name;
            traffic-statistics {
                file filename <world-readable | no-world-readable>;
                interval seconds;
            }
        }
    }
}
```



```

    }
  }
  route-target {
    accepted-prefix-limit {
      maximum number;
      proxy-generate <route-target-policy route-target-policy-name>;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
  }
}
(inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
  signaling {
    accepted-prefix-limit {
      maximum number;
      teardown <percentage-threshold> idle-timeout (forever | minutes);
    }
    add-path {
      send {
        path-count number;
        prefix-policy [ policy-names ];
      }
      receive;
    }
    aigp [disable];
    damping;
    loops number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
  }
}
}

```

Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <i>neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp group <i>group-name</i> <i>neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>inet-mvpn and inet6-mvpn statements introduced in Junos OS Release 8.4.</p> <p>inet-mdt statement introduced in Junos OS Release 9.4.</p> <p>Support for the loops statement introduced in Junos OS Release 9.6.</p>
Description	<p>Enable multiprotocol BGP (MP-BGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, to specify MP-BGP to carry NLRI for the IPv6 address family, or to carry NLRI for VPNs.</p>

- Options**
- any**—Configure the family type to be both unicast and multicast.
 - inet**—Configure NLRI parameters for IPv4.
 - inet6**—Configure NLRI parameters for IPv6.
 - inet-mdt**—Configure NLRI parameters for the multicast distribution tree (MDT) subaddress family identifier (SAFI) for IPv4 traffic in Layer 3 VPNs.
 - inet-mvpn**—Configure NLRI parameters for IPv4 for multicast VPNs.
 - inet6-mvpn**—Configure NLRI parameters for IPv6 for multicast VPNs.
 - inet-vpn**—Configure NLRI parameters for IPv4 for Layer 3 VPNs.
 - inet6-vpn**—Configure NLRI parameters for IPv6 for Layer 3 VPNs.
 - iso-vpn**—Configure NLRI parameters for IS-IS for Layer 3 VPNs.
 - l2vpn**—Configure NLRI parameters for IPv4 for MPLS-based Layer 2 VPNs and VPLS.
 - labeled-unicast**—Configure the family type to be labeled-unicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by labeled-unicast for resolving the labeled-unicast routes. This statement is supported only with **inet** and **inet6**.
 - multicast**—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.
 - unicast**—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes. The default family type is **unicast**.

The remaining statements are explained separately.

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

- Related Documentation**
- [autonomous-system on page 94](#)
 - [local-as on page 229](#)
 - Understanding Multiprotocol BGP

fate-sharing

Syntax	<pre>fate-sharing { group <i>group-name</i> { cost <i>value</i>; from <i>address</i> <to <i>address</i>>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Specify a backup path in case the primary path becomes unusable.</p> <p>You specify one or more objects with common characteristics within a group. All objects are treated as /32 host addresses. The objects can be a LAN interface, a router ID, or a point-to-point link. Sequence is insignificant.</p> <p>Changing the fate-sharing database does not affect existing established LSPs until the next CSPF reoptimization. The fate-sharing database does affect fast-reroute detour path computations.</p>
Options	<p>cost <i>value</i>—Cost assigned to the group. Range: 1 through 65,535 Default: 1</p> <p>from <i>address</i>—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate from addresses in the group.</p> <p>group <i>group-name</i>—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.</p> <p>to <i>address</i>—(Optional) Address of egress router. For point-to-point link objects, you must specify both a from and a to address.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring the Ingress Router for MPLS-Signaled LSPsJunos OS MPLS Applications Configuration Guide

flow

Syntax	<pre> flow { route <i>name</i> { match { <i>match-conditions</i>; } term-order (legacy standard); then { <i>actions</i>; } } validation { traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } } } </pre>
Hierarchy Level	[edit routing-options]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>term-order statement introduced in Junos OS Release 10.0</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure a flow route.
Default	legacy
Options	<p>actions—An action to take if conditions match.</p> <p>match-conditions—Match packets to these conditions.</p> <p>route name—Name of the flow route.</p> <p>standard—Specify to use version 7 or later of the flow-specification algorithm.</p> <p>term-order (legacy standard)—Specify the version of the flow-specification algorithm.</p> <ul style="list-style-type: none"> legacy—Use version 6 of the flow-specification algorithm. standard—Use version 7 of the flow-specification algorithm. <p>then—Actions to take on matching packets.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- Example: Configuring Flow Routes

flow-map

Syntax	<pre>flow-map <i>flow-map-name</i> { bandwidth (<i>bps</i> <i>adaptive</i>); forwarding-cache { timeout (<i>never non-discard-entry-only</i> <i>minutes</i>); } policy [<i>policy-names</i>]; redundant-sources [<i>addresses</i>]; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure multicast flow maps.
Options	<p><i>flow-map-name</i>—Name of the flow-map.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring a Multicast Flow Map

forwarding-cache (Flow Maps)

Syntax	<pre>forwarding-cache { timeout (minutes never non-discard-entry-only); }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-options multicast flow-map <i>flow-map-name</i>]</pre>
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure multicast forwarding cache properties for the flow map.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

forwarding-cache (Multicast)

Syntax	<pre>forwarding-cache { family (inet inet6) { threshold { log-warning value; suppress value <reuse value>; } } threshold { log-warning value; suppress value <reuse value>; } timeout minutes; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure multicast forwarding cache properties. These properties include threshold suppression and reuse limits, the threshold at which a warning message is logged, and timeout values.</p> <p>Specify a value for the threshold at which to suppress new multicast forwarding cache entries and an optional reuse value for the threshold at which the router begins to create new multicast forwarding cache entries. The range for both is from 1 through 200,000. If configured, the reuse value should be less than the suppression threshold value. The suppression value is mandatory. If you do not specify the optional reuse value, then the number of multicast forwarding cache entries is limited to the suppression value. A new entry is created as soon as the number of multicast forwarding cache entries falls below the suppression value.</p> <p>You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the threshold statement globally for the multicast forwarding cache or including the family statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p>
Default	By default, there are no limits on the number of multicast forwarding cache entries.
Options	<p>family (inet inet6)—(Optional) Apply the configured thresholds to either IPv4 or IPv6 multicast forwarding cache entries.</p> <p>Default: By default, the configured thresholds are applied to both IPv4 and IPv6 multicast forwarding cache entries.</p>

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring the Multicast Forwarding Cache

forwarding-table

Syntax	<pre>forwarding-table { chained-composite-next-hop { ingress { l3vpn { extended-space; } } } export [<i>policy-name</i>]; indexed-next-hop; (indirect-next-hop no-indirect-next-hop); (indirect-next-hop-change-acknowledgements no-indirect-next-hop-change-acknowledgements); krt-nexthop-ack-timeout <i>interval</i>; unicast-reverse-path (active-paths feasible-paths); }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure information about the routing device's forwarding table. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Load Balancing BGP Traffic

generate

Syntax	<pre>generate { defaults { generate-options; } route destination-prefix { policy policy-name; generate-options; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>], [edit routing-options], [edit routing-options rib <i>routing-table-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure generated routes, which are used as routes of last resort.
Options	<p>defaults—(Optional) Specify global generated route options. These options only set default attributes inherited by all newly created generated routes. These are treated as global defaults and apply to all the generated routes you configure in the generate statement.</p> <p>generate-options—Additional information about generated routes, which is included with the route when it is installed in the routing table. Specify zero or more of the following options in generate-options. Each option is explained separately.</p> <ul style="list-style-type: none">• (active passive);• as-path <i><as-path></i> <i><origin (egp igp incomplete)></i> <i><atomic-aggregate></i> <i><aggregator as-number in-address></i>;• (brief full);• community [<i>community-ids</i>];• discard;• (metric <i>metric2</i> <i>metric3</i> <i>metric4</i>) <i>value</i> <i><type type></i>;• (preference <i>preference2</i> color <i>color2</i>) <i>preference</i> <i><type type></i>;• tag <i>string</i>; <p>route destination-prefix—Configure a non-default generated route:</p> <ul style="list-style-type: none">• default—For the default route to the destination. This is equivalent to specifying an IP address of 0.0.0.0/0.

- *destination-prefix/prefix-length—/destination-prefix* is the network portion of the IP address, and *prefix-length* is the destination prefix length.

The [policy](#) 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	<ul style="list-style-type: none">• Example: Conditionally Generating Static Routes

graceful-restart (Protocols BGP)

Syntax	<pre>graceful-restart { disable; restart-time <i>seconds</i>; stale-routes-time <i>seconds</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>]</pre>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Enable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the restart-time statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the stale-routes-time statement at the [edit protocols bgp graceful-restart] hierarchy level.</p>



NOTE: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.

For graceful restart to function properly, graceful restart must be enabled at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level as well as in the protocol level.

For example:

```
protocols {
  bgp {
    group ext {
      graceful-restart;
    }
  }
}
routing-options {
  graceful-restart;
}
```

Graceful restart is enabled both at the **[edit routing-options]** hierarchy level, as well as at the routing protocol level. If graceful restart is not configured in both sections, the peer might have its route removed after a restart, which is not the intended behavior.

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	• Configuring Graceful Restart Options for BGP
	• Configuring Graceful Restart for QFabric Systems
	• Junos OS High Availability Configuration Guide

graceful-restart (Protocols IS-IS)

Syntax	<pre>graceful-restart { disable; helper-disable; restart-duration <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure graceful restart parameters for IS-IS.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level. When graceful restart is enabled, the restarting routing device is not removed from the network topology during the restart period. The adjacencies are reestablished after restart is complete.</p> <p>On LAN interfaces where IS-IS is configured on a transit router that serves as the designated router (DR), a graceful restart causes:</p> <ul style="list-style-type: none">• The ingress router of the label-switched path (LSP), which passes through the DR, to break the LSP.• The ingress router to re-signal the LSP.
Options	<p>disable—Disable graceful restart for IS-IS.</p> <p>helper-disable—Disable graceful restart helper capability. Helper mode is enabled by default.</p> <p>restart-duration <i>seconds</i>—Time period for the restart to last, in seconds. Range: 30 through 300 seconds Default: 30 seconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Routing Protocols Graceful Restart

graceful-restart (Protocols OSPF)

Syntax	<pre> graceful-restart { disable; helper-disable (standard restart-signaling both); no-strict-lsa-checking; notify-duration <i>seconds</i>; restart-duration <i>seconds</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the no-strict-lsa-checking statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode standard, restart-signaling, and both options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure graceful restart for OSPF.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level.</p>
Options	<p>disable—Disable graceful restart for OSPF.</p> <p>helper-disable (standard restart-signaling both)—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The standard, restart-signaling, and both options are only supported for OSPFv2. Specify standard to disable helper mode for standard graceful restart (based on RFC 3623). Specify restart-signaling to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify both to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p>Default: Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p>no-strict-lsa-checking—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p>



NOTE: The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both

statements at the same time, the routing device displays a warning message when you enter the `show protocols (ospf | ospf3)` command.

.....
notify-duration seconds—Estimated time needed to send out purged grace LSAs over all the interfaces.

Range: 1 through 3600 seconds

Default: 30 seconds

restart-duration seconds—Estimated time needed to reacquire a full OSPF neighbor from each area.

Range: 1 through 3600 seconds

Default: 180 seconds

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

Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Graceful Restart for OSPF • Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart • Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart • Example: Disabling Strict LSA Checking for OSPF Graceful Restart • Configuring Graceful Restart for QFabric Systems • Junos OS High Availability Configuration Guide
------------------------------	--

graceful-restart (Protocols RIP)

Syntax	<pre>graceful-restart { disable; restart-time <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rip], [edit protocols rip]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure graceful restart for RIP.
Options	disable —Disables graceful restart for RIP. 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	<ul style="list-style-type: none">Junos OS High Availability Configuration Guide

graceful-restart (Protocols RIPng)

Syntax	<pre>graceful-restart { disable; restart-time <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ripng], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng], [edit protocols ripng], [edit routing-instances <i>routing-instance-name</i> protocols ripng]
Release Information	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.
Description	Configure graceful restart for RIPng.
Options	disable —Disables graceful restart for RIPng. 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	<ul style="list-style-type: none">Junos OS High Availability Configuration Guide

graceful-restart (Enabling Globally)

Syntax	<pre> graceful-restart { disable; helper-disable; maximum-helper-recovery-time <i>seconds</i>; maximum-helper-restart-time <i>seconds</i>; notify-duration <i>seconds</i>; recovery-time <i>seconds</i>; restart-duration <i>seconds</i>; stale-routes-time <i>seconds</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally.</p> <p>For VPNs, the graceful-restart statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.</p> <p>For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p>
Default	Graceful restart is disabled by default.
Options	The remaining statements are explained separately.
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Enabling Graceful Restart Configuring Routing Protocols Graceful Restart Configuring Graceful Restart for MPLS-Related Protocols Configuring VPN Graceful Restart Configuring Logical System Graceful Restart Graceful Restart Configuration Statements Configuring Graceful Restart for QFabric Systems

group (Protocols BGP)

```
Syntax  group group-name {
    advertise-inactive;
    allow [ network/mask-length ];
    authentication-key key;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family {
        (inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {
            (any | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                add-path {
                    send {
                        path-count number;
                        prefix-policy [ policy-names ];
                    }
                    receive;
                }
                aigp [disable];
                damping;
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                rib-group group-name;
                topology name {
                    community {
                        target identifier;
                    }
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib inet.3;
        }
    }
}
```

```

        rib-group group-name;
    }
}
route-target {
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-preference local-preference;
log-updown;
metric-out metric;
multihop <ttl-value>;
multipath {
    multiple-as;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
type type;
neighbor address {
    ... peer-specific-options ...
}
}

```

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

Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Define a BGP peer group. BGP peer groups share a common type, peer autonomous system (AS) number, and cluster ID, if present. To configure multiple BGP groups, include multiple group statements.</p> <p>By default, the group's options are identical to the global BGP options. To override the global options, include group-specific options within the group statement.</p> <p>The group statement is one of the statements you must include in the configuration to run BGP on the routing device.</p> <p>Each group must contain at least one peer.</p>
Options	<p>group-name—Name of the BGP group.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• BGP Configuration Guide

group (Protocols RIP)

```

Syntax  group group-name {
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            transmit-interval {
                threshold milliseconds;
                minimum-interval milliseconds;
            }
            multiplier number;
            version (0 | 1 | automatic);
        }
        demand-circuit;
        export policy;
        max-retrans-time seconds;
        metric-out metric;
        preference number;
        route-timeout seconds;
        update-interval seconds;
        neighbor neighbor-name {
            authentication-key password;
            authentication-type type;
            bfd-liveness-detection {
                authentication {
                    algorithm algorithm-name;
                    key-chain key-chain-name;
                    loose-check;
                }
                detection-time {
                    threshold milliseconds;
                }
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                transmit-interval {
                    threshold milliseconds;
                    minimum-interval milliseconds;
                }
                multiplier number;
                version (0 | 1 | automatic);
            }
            (check-zero | no-check-zero);
            demand-circuit;
            import policy-name;
            max-retrans-time seconds;
            message-size number;

```

```
metric-in metric;  
metric-out metric;  
receive receive-options;  
route-timeout seconds;  
send send-options;  
update-interval seconds;  
}  
}
```

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure a set of RIP neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group. Each group must contain at least one neighbor. You should create a group for every export policy.

Options *group-name*—Name of a group, up to 16 characters long.

The remaining statements are explained separately.

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

Related Documentation

- Example: Configuring RIP


group (Protocols RIPng)

Syntax	<pre> group <i>group-name</i> { export [<i>policy-names</i>]; metric-out <i>metric</i>; neighbor <i>neighbor-name</i> { import <i>policy-name</i>; metric-in <i>metric</i>; receive <none>; route-timeout <i>seconds</i>; send <none>; update-interval <i>seconds</i>; } preference <i>number</i>; route-timeout <i>seconds</i>; update-interval <i>seconds</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ripng],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng],</p> <p>[edit protocols ripng],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>
Description	<p>Configure a set of RIPng neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group.</p> <p>Each group must contain at least one neighbor. You should create a group for each export policy that you have.</p>
Options	<p><i>group-name</i>—Name of a group, up to 16 characters long.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring RIPng

hello-authentication-key

Syntax	hello-authentication-key <i>password</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>], [edit protocols isis interface <i>interface-name</i> level <i>number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure an authentication key (password) for hello packets. Neighboring routing devices use the password to verify the authenticity of packets sent from an interface. For the key to work, you also must include the hello-authentication-type statement.
Default	By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.
Options	password —Authentication password. The password can be up to 255 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• authentication-key on page 88• authentication-type on page 92• hello-authentication-type on page 181

hello-authentication-type

Syntax	hello-authentication-type (md5 simple);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Enable authentication on an interface for hello packets. If you enable authentication on hello packets, you must specify a password by including the hello-authentication-key statement.</p> <p>You can configure authentication for a given IS-IS level on an interface. On a point-to-point link, if you enable hello authentication for both IS-IS levels, the password configured for Level 1 is used for both levels.</p>
	<div>  <p>CAUTION: If no authentication is configured for Level 1 on a point-to-point link with both levels enabled, the hello packets are sent without any password, regardless of the Level 2 authentication configurations.</p> </div>
Default	By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.
Options	<p>md5—Specifies Message Digest 5 as the packet verification type.</p> <p>simple—Specifies simple authentication as the packet verification type.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • authentication-key on page 88 • authentication-type on page 92 • hello-authentication-key on page 180

hello-interval (Protocols IS-IS)

Syntax	<code>hello-interval <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Modify the frequency with which the routing device sends hello packets out of an interface, in seconds.</p> <p>Routing devices send hello packets at a fixed interval on all interfaces to establish and maintain neighbor relationships. This interval is advertised in the hello interval field in the hello packet.</p> <p>You can send out hello packets in subsecond intervals. To send out hello packets every 333 milliseconds, set the hello-interval value to 1.</p>
Options	<i>seconds</i> —Frequency of transmission for hello packets. Range: 1 through 20,000 seconds Default: 3 seconds (for designated intermediate system [DIS] routers), 9 seconds (for non-DIS routers)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">hold-time

hello-interval (Protocols OSPF)

Syntax	<code>hello-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</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) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Specify how often the routing device sends hello packets out the interface. The hello interval must be the same for all routing devices on a shared logical IP network.
Options	<p>seconds—Time between hello packets, in seconds.</p> <p>Range: 1 through 255 seconds</p> <p>Default: 10 seconds (broadcast and point-to-point networks); 30 seconds (nonbroadcast multiple access [NBMA] networks)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring OSPF Timers • Configuring RSVP and OSPF for LMP Peer Interfaces • dead-interval on page 133

hello-padding

Syntax	hello-padding (adaptive disable loose strict);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure padding on hello packets to accommodate asymmetrical maximum transfer units (MTUs) from different hosts.</p> <p>This helps to prevent a premature adjacency Up state when one routing device's MTU does not meet the requirements to establish the adjacency.</p> <p>As an OSI Layer 2 protocol, IS-IS does not support data fragmentation. Therefore, maximum packet sizes must be established and supported between two routers. During adjacency establishment, the IS-IS protocol makes sure that the link supports a packet size of 1492 bytes by padding outgoing hello packets up to the maximum packet size of 1492 bytes.</p> <p>This is the default behavior of the Junos OS IS-IS implementation. However, Junos OS provides an option to disable hello padding that can override this behavior.</p> <p>There are four types of hello padding:</p> <ul style="list-style-type: none">• Adaptive padding—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state type, length, and value (TLV) tuple. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. Adaptive padding has more overhead than loose padding and is able to detect MTU asymmetry from one side of the connection. This one-sided detection can result in generation of extra link-state PDUs that are flooded throughout the network. Specify the adaptive option to configure enough padding to establish an adjacency to neighbors.• Disabled padding—Padding is disabled on all types of interfaces for all adjacency states. Specify the disable option to accommodate interfaces that support less than the default packet size of 1492 bytes.• Loose padding (the default)—The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. Loose padding might not be able to detect certain situations such as asymmetrical MTUs between the routing devices. Specify the loose option to configure enough padding to initialize an adjacency to neighbors.

- Strict padding—Padding is done on all interface types and for all adjacency states, and is continuous. Strict padding has the most overhead. The advantage is that strict padding detects MTU issues on both sides of a link. Specify the **strict** option to configure padding to allow all adjacency states with neighbors.

Options **adaptive**—Configure padding until the neighbor adjacency is established and active.

disable—Disable padding on all types of interfaces for all adjacency states.

loose—Configure padding until the state of the adjacency is initialized.

strict—Configure padding for all adjacency states.

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

Related Documentation

- Example: Configuring IS-IS

holddown (Protocols RIP)

Syntax `holddown seconds;`

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure how long the expired route is retained in the routing table before being removed.

When the hold-down timer runs on RIP demand circuits, routes are advertised as unreachable on other interfaces. When the hold-down timer expires, the route is removed from the routing table if all destinations detect that the route is unreachable or the remaining destinations are down.

Options **seconds**—Estimated time to wait before making updates to the routing table.

Range: 10 through 180 seconds

Default: 180 seconds

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

Related Documentation

- Example: Configuring RIP Timers
- RIP Demand Circuits Overview

holddown (Protocols RIPng)

Syntax	<code>holddown seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ripng], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng], [edit protocols ripng], [edit routing-instances <i>routing-instance-name</i> protocols ripng]
Release Information	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.
Description	Configure how long the expired route is retained in the routing table before being removed.
Options	seconds —Estimated time to wait before removing expired routes from the routing table. Default: 180 seconds Range: 10 through 180 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring RIPng Timers

hold-time

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	Configure the time in seconds after which a backup router with the highest priority preempts the master router.
Options	seconds —Hold-time period.
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">Configuring VRRP for IPv6 (CLI Procedure)

hold-time (Protocols BGP)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the hold-time value to use when negotiating a connection with the peer. The hold-time value is advertised in open packets and indicates to the peer the length of time that it should consider the sender valid. If the peer does not receive a keepalive, update, or notification message within the specified hold time, the BGP connection to the peer is closed and routing devices through that peer become unavailable.</p> <p>The hold time is three times the interval at which keepalive messages are sent.</p> <p>BGP on the local routing device uses the smaller of either the local hold-time value or the peer's hold-time value received in the open message as the hold time for the BGP connection between the two peers.</p> <p>Starting in Junos OS Release 12.3, the BGP hold-time value can be zero (0). This implies that the speaker does not expect keepalive messages from its peer to maintain the BGP session. When negotiating between two peers, if one side requests a nonzero hold time and the other requests a zero hold time, the negotiation settles on the nonzero value and keepalive intervals are determined accordingly. Both sides must be set to zero for keepalive messages to stop being sent.</p>
Options	<p>seconds—Hold time.</p> <p>Range: 10 through 65,535 seconds (or 0 for infinite hold time)</p> <p>Default: 90 seconds</p>



TIP: When you set a hold-time value of 1 through 19 seconds, we recommend that you also configure the BGP `precision-timers` statement. The `precision-timers` statement ensures that if scheduler slip messages occur, the routing device continues to send keepalive messages. When the `precision-timers` statement is included, keepalive message generation is performed in a dedicated kernel thread, which helps to prevent BGP session flaps.

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

Related Documentation

- [BGP Messages Overview](#)
- [precision-timers](#)

hold-time (Protocols IS-IS)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Set the length of time a neighbor considers this router to be operative (up) after receiving a hello packet. If the neighbor does not receive another hello packet within the specified time, it marks this routing device as inoperative (down). The hold time itself is advertised in the hello packets.</p> <p>The hold time specifies how long a neighbor should consider this routing device to be operative without receiving another hello packet. If the neighbor does not receive a hello packet from this routing device within the hold time, it marks the routing device as being unavailable.</p>
Options	<p>seconds—Hold-time value, in seconds.</p> <p>Range: 3 through 65,535 seconds, or 1 to send out hello packets every 333 milliseconds</p> <p>Default: 9 seconds (for designated intermediate system [DIS] routers), 27 seconds (for non-DIS routers; three times the default hello interval)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring IS-IS hello-interval on page 182

idle-after-switch-over

Syntax	<code>idle-after-switch-over (forever seconds);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the routing device so that it does not automatically reestablish BGP peer sessions after a nonstop active routing (NSR) switchover. This feature is particularly useful if you are using dynamic routing policies because the dynamic database is not synchronized with the backup Routing Engine when NSR is enabled.
Options	forever —Do not reestablish a BGP peer session after an non-stop routing switchover until the clear bgp neighbor command is issued. seconds —Do not reestablish a BGP peer session after an non-stop routing switchover until after the specified period. Range: 1 through 4,294,967,295 ($2^{32} - 1$)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Preventing Automatic Reestablishment of BGP Peer Sessions After NSR SwitchoversRouting Policy Configuration GuideJunos OS High Availability Configuration Guide

ignore-attached-bit

Syntax	ignore-attached-bit;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Ignore the attached bit on IS-IS Level 1 routers. Configuring this statement enables the routing device to ignore the attached bit on incoming Level 1 link-state PDUs. If the attached bit is ignored, no default route, which points to the routing device which has set the attached bit, is installed.</p> <p>There might be times, such as during a denial-of-service (DoS) attack, that you do not want a Level 1 router to be able to forward traffic based on a default route.</p> <p>To prevent a routing device from being able to reach interarea destinations, you can prevent the routing device from installing the default route without affecting the status of its IS-IS adjacencies. The ignore-attached-bit statement is used to tell the routing device to ignore the presence of the attached bit in Level 1 link-state PDUs, which blocks the installation of the IS-IS default route.</p>
Default	The ignore-attached-bit statement is disabled by default.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> •

ignore-lsp-metrics

Syntax	ignore-lsp-metrics;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf traffic-engineering shortcuts], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf traffic-engineering shortcuts], [edit protocols ospf traffic-engineering], [edit routing-instances <i>routing-instance-name</i> protocols ospf traffic-engineering shortcuts]
Release Information	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for (OSPFv3) introduced in Junos OS Release 9.4. Support for (OSPFv3) introduced in Junos OS Release 9.4 for EX Series switches.
Description	Ignore RSVP LSP metrics in OSPF traffic engineering shortcut calculations.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Enabling OSPF Traffic Engineering Support

import (Protocols BGP)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor 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 <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name</i> <i>neighbor 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 <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Apply one or more routing policies to routes being imported into the Junos OS routing table from BGP.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BGP Interactions with IGPs • Understanding Route Advertisement • Importing and Exporting Routes • Routing Policy Configuration Guide • export on page 145

import (Protocols OSPF)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> (ospf ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit protocols (ospf ospf3)],</code> <code>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Filter OSPF routes from being added to the routing table.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding OSPF Routing Policy• Import and Export Policies for Network Summaries Overview• export on page 147• Routing Policy Configuration Guide

import (Protocols RIP)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rip],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip 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 rip],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols rip],</p> <p>[edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Apply one or more policies to routes being imported by the local routing device from neighbors.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Applying Policies to RIP Routes Imported from Neighbors • Routing Policy Configuration Guide • export on page 148


import (Protocols RIPng)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols ripng],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ripng 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> ripng],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ripng group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit protocols ripng],</code> <code>[edit protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ripng],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> neighbor</code> <code> <i>neighbor-name</i>]</code>
Release Information	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.
Description	Apply one or more policies to routes being imported into the local routing device from its neighbors.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Applying Policies to RIPng Routes Imported from Neighbors• export on page 149


import (Routing Options)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution rib],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options resolution rib],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options resolution rib],</p> <p>[edit routing-options resolution rib]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Specify one or more import policies to use for route resolution.
Options	<i>policy-names</i> —Name of one or more import policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Route Resolution on PE Routers

import-policy

Syntax	<code>import-policy [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options rib-groups <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>], [edit routing-options rib-groups <i>group-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Apply one or more policies to routes imported into the routing table group. The import-policy statement complements the import-rib statement and cannot be used unless you first specify the routing tables to which routes are being imported.
<div> NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.</div>	
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Exporting Specific Routes from One Routing Table Into Another Routing Table• export-rib on page 151• passive


import-rib

Syntax	<code>import-rib [<i>routing-table-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib-groups <i>group-name</i>],</p> <p>[edit routing-options rib-groups <i>group-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the name of the routing table into which Junos OS should import routing information. The first routing table name you enter is the primary routing table. Any additional names you enter identify secondary routing tables. When a protocol imports routes, it imports them into the primary and any secondary routing tables. If the primary route is deleted, the secondary route also is deleted. For IPv4 import routing tables, the primary routing table must be inet.0 or routing-instance-name.inet.0. For IPv6 import routing tables, the primary routing table must be inet6.0.</p> <p>In Junos OS Release 9.5 and later, you can configure an IPv4 import routing table that includes both IPv4 and IPv6 routing tables. Including both types of routing tables permits you, for example, to populate an IPv6 routing table with IPv6 addresses that are compatible with IPv4. In releases prior to Junos OS Release 9.5, you could configure an import routing table with only either IPv4 or IPv6 routing tables.</p>
	<div>  <p>NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.</p> </div>
Options	<i>routing-table-names</i> —Name of one or more routing tables.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Exporting Specific Routes from One Routing Table Into Another Routing Table • export-rib on page 151 • passive

include-mp-next-hop

Syntax	include-mp-next-hop;
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>]</pre>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable multiprotocol updates to contain next-hop reachability information.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Examples: Configuring Multiprotocol BGP

indirect-next-hop

Syntax	(indirect-next-hop no-indirect-next-hop);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options forwarding-table], [edit routing-options forwarding-table]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable indirectly connected next hops for route convergence. This statement is implemented on the Packet Forward Engine to speed up forwarding information base (FIB) updates. Configuring this statement significantly speeds convergence times. The only downside of configuring this statement is that some additional FIB memory overhead is required. Unless routes have an extremely high number of next hops, this increased memory usage should not be noticeable.
<div>  <p>NOTE:</p> <ul style="list-style-type: none"> When virtual private LAN service (VPLS) is configured on the routing device, the <code>indirect-next-hop</code> statement is configurable at the [edit routing-options <code>forwarding-table</code>] hierarchy level. However, this configuration is not applicable to indirect nexthops specific to VPLS routing instances. By default, the Junos Trio Modular Port Concentrator (MPC) chipset on MX Series routers is enabled with indirectly connected next hops, and this cannot be disabled using the <code>no-indirect-next-hop</code> statement. </div>	
Default	Disabled.
Options	<code>indirect-next-hop</code> —Enable indirectly connected next hops. <code>no-indirect-next-hop</code> —Explicitly disable indirect next hops.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine

install (Routing Options)

Syntax	(install no-install);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)] [edit routing-options static (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure whether Junos OS installs all static routes into the forwarding table. Even if you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols.
Options	install —Explicitly install all static routes into the forwarding table. Include this statement when configuring an individual route in the route portion of the static statement to override a no-install option specified in the defaults portion of the statement. no-install —Do not install the route into the forwarding table, even if it is the route with the lowest preference. Default: install
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Examples: Configuring Static Routesstatic on page 376

instance-export

Syntax	<code>instance-export [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Apply one or more policies to routes being exported from a routing instance.
Options	<i>policy-names</i> —Name of one or more export policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring and Verifying the Auto Export Feature • Routing Policy Configuration Guide

instance-import

Syntax	<code>instance-import [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Apply one or more policies to routes being imported into a routing instance.
Options	<i>policy-names</i> —Name of one or more import policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring and Verifying the Auto Export Feature • Routing Policy Configuration Guide

inter-area-prefix-export

Syntax	<code>inter-area-prefix-export [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <i>area</i> <i>area-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <i>area</i> <i>area-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ip4-unicast ipv4-multicast ipv6-multicast) <i>area</i> <i>area-id</i>],</code> <code>[edit protocols ospf3 <i>area</i> <i>area-id</i>],</code> <code>[edit protocols ospf3 <i>realm</i> (ip4-unicast ipv4-multicast ipv6-multicast) <i>area</i> <i>area-id</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>area</i> <i>area-id</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ip4-unicast ipv4-multicast ipv6-multicast) <i>area</i> <i>area-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.
Description	Apply an export policy for OSPFv3 to specify which interarea prefix link-state advertisements (LSAs) are flooded into an area.
Options	<i>policy-name</i> —Name of a policy configured at the <code>[edit policy-options policy-statement <i>policy-name</i> term <i>term-name</i>]</code> hierarchy level.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Import and Export Policies for Network Summaries Overview• inter-area-prefix-import on page 205• Routing Policy Configuration Guide

inter-area-prefix-import

Syntax	<code>inter-area-prefix-import [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 area <i>area-id</i>],</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) area <i>area-id</i>],</p> <p>[edit protocols ospf3 area <i>area-id</i>],</p> <p>[edit protocols ospf3 realm (ip4-unicast ipv4-multicast ipv6-multicast)], area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Apply an import policy for OSPFv3 to specify which routes learned from an area are used to generate interarea prefixes into other areas.
Options	<i>policy-name</i> —Name of a policy configured at the [edit policy-options policy-statement <i>policy-name</i> term <i>term-name</i>] hierarchy level.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Import and Export Policies for Network Summaries Overview inter-area-prefix-export on page 204 Routing Policy Configuration Guide

interface (Protocols IS-IS)

```
Syntax  interface (all | interface-name) {
        disable;
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            transmit-interval {
                threshold milliseconds;
                minimum-interval milliseconds;
            }
            multiplier number;
        }
        checksum;
        csnp-interval (seconds | disable);
        hello-padding (adaptive | loose | strict);
        ldp-synchronization {
            disable;
            hold-time seconds;
        }
        lsp-interval milliseconds;
        mesh-group (value | blocked);
        no-adjacency-holddown;
        no-ipv4-multicast;
        no-ipv6-multicast;
        no-ipv6-unicast;
        no-unicast-topology;
        passive;
        point-to-point;
        level level-number {
            disable;
            hello-authentication-key key;
            hello-authentication-key-chain key-chain-name;
            hello-authentication-type authentication;
            hello-interval seconds;
            hold-time seconds;
            ipv4-multicast-metric metric;
            ipv6-multicast-metric metric;
            ipv6-unicast-metric metric;
            metric metric;
            passive;
            priority number;
            te-metric metric;
        }
    }
```

Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure interface-specific IS-IS properties. To configure more than one interface, include the interface statement multiple times.</p> <p>Enabling IS-IS on an interface (by including the interface statement at the [edit protocols isis] or the [edit routing-instances <i>routing-instance-name</i> protocols isis] hierarchy level), disabling it (by including the disable statement), and not actually having IS-IS run on an interface (by including the passive statement) are mutually exclusive states.</p>
Options	<p>all—Have Junos OS create IS-IS interfaces automatically. If you include this option, disable IS-IS on the management interface (fxp0).</p> <p>interface-name—Name of an interface. Specify the full interface name, including the physical and logical address components.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS • Example: Configuring Multi-Level IS-IS

interface (Protocols OSPF)

Syntax interface *interface-name* {
 disable;
 authentication *key* <*key-id identifier*>;
 bfd-liveness-detection {
 authentication {
 algorithm *algorithm-name*;
 key-chain *key-chain-name*;
 loose-check;
 }
 detection-time {
 threshold *milliseconds*;
 }
 minimum-interval *milliseconds*;
 minimum-receive-interval *milliseconds*;
 transmit-interval {
 threshold *milliseconds*;
 minimum-interval *milliseconds*;
 }
 multiplier *number*;
 }
 dead-interval *seconds*;
 demand-circuit;
 hello-interval *seconds*;
 ipsec-sa *name*;
 interface-type *type*;
 ldp-synchronization {
 disable;
 hold-time *seconds*;
 }
 metric *metric*;
 neighbor *address* <*eligible*>;
 no-interface-state-traps;
 passive;
 poll-interval *seconds*;
 priority *number*;
 retransmit-interval *seconds*;
 te-metric *metric*;
 topology (ipv4-multicast | *name*) {
 metric *metric*;
 }
 transit-delay *seconds*;
 transmit-interval *seconds*;
}

Hierarchy Level [edit logical-systems *logical-system-name* protocols (ospf | ospf3) *area area-id*],
 [edit logical-systems *logical-system-name* protocols ospf3 *realm* (ipv4-unicast |
 ipv4-multicast | ipv6-multicast) *area area-id*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 (ospf | ospf3) *area area-id*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 ospf3 *realm* (ipv4-unicast | ipv4-multicast | ipv6-multicast) *area area-id*],
 [edit protocols (ospf | ospf3) *area area-id*],

```
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast |
  ipv4-multicast | ipv6-multicast) area area-id]
```

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Support for the **topology** statement introduced in Junos OS Release 9.0.
 Support for the **topology** statement introduced in Junos OS Release 9.0 for EX Series switches.
 Support for the **realm** statement introduced in Junos OS Release 9.2.
 Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.
 Support for the **no-interface-state-traps** statement introduced in Junos OS Release 10.3.
 This statement is supported only for OSPFv2.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Enable OSPF routing on a routing device interface.

You must include at least one **interface** statement in the configuration to enable OSPF on the routing device.

Options **interface-name**—Name of the interface. Specify the interface by IP address or interface name for OSPFv2, or only the interface name for OSPFv3. Using both the interface name and IP address of the same interface produces an invalid configuration. To configure all interfaces, you can specify **all**. Specifying a particular interface and **all** produces an invalid configuration.



NOTE: For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as **interface-name**.

The remaining statements are explained separately.



NOTE: You cannot run both OSPF and **ethernet-tcc** encapsulation between two Juniper Networks routing devices.

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

- Related Documentation**
- OSPF Configuration Overview
 - Example: Configuring Multitopology Routing Based on Applications
 - Example: Configuring Multitopology Routing Based on a Multicast Source
 - Example: Configuring Multiple Address Families for OSPFv3
 - neighbor

interface (Routing Options)

Syntax

```
interface interface-names {  
    maximum-bandwidth bps;  
    no-qos-adjust;  
    reverse-oif-mapping {  
        no-qos-adjust;  
    }  
    subscriber-leave-timer seconds;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name  
    routing-options multicast],  
[edit logical-systems logical-system-name routing-options multicast],  
[edit routing-instances routing-instance-name routing-options multicast],  
[edit routing-options multicast]
```

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Enable multicast traffic on an interface.



TIP: You cannot enable multicast traffic on an interface by using the `routing-options multicast interface` statement and configure PIM on the interface.

Options

interface-name—Names of the physical or logical interface.

The remaining statements are explained separately.

Required Privilege Level

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

- Related Documentation**
- Example: Defining Interface Bandwidth Maximums
 - Example: Configuring Multicast with Subscriber VLANs

interface (Multicast Static Routes)

Syntax	<pre> interface <i>interface-names</i> { disable; maximum-bandwidth <i>bps</i>; no-qos-adjust; reverse-oif-mapping { no-qos-adjust; } subscriber-leave-timer <i>seconds</i>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast] </pre>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Enable multicast traffic on an interface.</p> <p>By default, multicast packets are forwarded by enabling Protocol Independent Multicast (PIM) on an interface. PIM adds multicast routes into the routing table.</p> <p>You can also configure multicast packets to be forwarded over a static route, such as a static route associated with an LSP next hop. Multicast packets are accepted on an interface and forwarded over a static route in the forwarding table. This is useful when you want to enable multicast traffic on a specific interface without configuring PIM on the interface.</p> <p>You cannot enable multicast traffic on an interface and configure PIM on the same interface simultaneously.</p> <p>Static routes must be configured before you can enable multicast on an interface. Configuring the interface statement alone does not install any routes into the routing table. This feature relies on the static route configuration.</p>
Options	<p><i>interface-names</i>—Name of one or more interfaces on which to enable multicast traffic.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Defining Interface Bandwidth Maximums Example: Configuring Multicast with Subscriber VLANs

interface-routes

Syntax

```
interface-routes {  
    family (inet | inet6) {  
        export {  
            lan;  
            point-to-point;  
        }  
    }  
    rib-group group-name;  
}
```

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.



NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Description Associate a routing table group with the routing device's interfaces, and specify routing table groups into which interface routes are imported.

By default, IPv4 interface routes (also called direct routes) are imported into routing table **inet.0**, and IPv6 interface routes are imported into routing table **inet6.0**. If you are configuring alternate routing tables for use by some routing protocols, it might be necessary to import the interface routes into the alternate routing tables. To define the routing tables into which interface routes are imported, you create a routing table group and associate it with the routing device's interfaces.

To create the routing table groups, include the **passive** statement at the **[edit routing-options]** hierarchy level.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group.

To export local routes, include the **export** statement.

To export LAN routes, include the **lan** option. To export point-to-point routes, include the **point-to-point** option.

Only local routes on point-to-point interfaces configured with a destination address are exportable.

Options **inet**—Specify the IPv4 address family.

inet6—Specify the IPv6 address family.

lan—Export LAN routes.

point-to-point—Export point-to-point routes.

 The remaining statements are explained separately.

Required Privilege routing—To view this statement in the configuration.

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


- Related** • Example: Importing Direct and Static Routes Into a Routing Instance
- Documentation** • Example: Configuring Multiple Routing Instances of OSPF
- passive

interface-type (Protocols OSPF)

Syntax	<code>interface-type (nbma p2mp p2p);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 for interface type p2p only introduced in Junos OS Release 9.4. You cannot configure other interface types for OSPFv3.</p> <p>Support for OSPFv3 for interface type p2p only introduced in Junos OS Release 9.4 for EX Series switches.</p>
Description	<p>Specify the type of interface.</p> <p>By default, the software chooses the correct interface type based on the type of physical interface. Therefore, you should never have to set the interface type. The exception to this is for NBMA interfaces, which default to an interface type of point-to-multipoint. To have these interfaces explicitly run in Nonbroadcast multiaccess (NBMA) mode, configure the nbma interface type, using the IP address of the local ATM interface.</p> <p>In Junos OS Release 9.3 and later, a point-to-point interface can be an Ethernet interface without a subnet.</p>
Default	The software chooses the correct interface type based on the type of physical interface.
Options	<p>nbma (OSPFv2 only)—Nonbroadcast multiaccess (NBMA) interface.</p> <p>p2mp (OSPFv2 only)—Point-to-multipoint interface.</p> <p>p2p—Point-to-point interface.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- About OSPF Interfaces
 - Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network

ipv4-multicast

Syntax	ipv4-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure alternate IPv4 multicast topologies.
	<div>  <p>NOTE: The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.</p> </div>
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS Multicast Topology

ipv4-multicast-metric

Syntax	ipv4-multicast-metric <i>metric</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the multicast topology metric value for the level.
Options	<i>metric</i> —Metric value. Range: 0 through 16,777,215
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring IS-IS Multicast Topology


ipv6-multicast

Syntax	ipv6-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure alternate IPv6 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring IS-IS Multicast Topology

ipv6-multicast-metric

Syntax	<code>ipv6-multicast-metric <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Specify the IPv6 alternate multicast topology metric value for the level.
Options	<p><i>metric</i>—Metric value.</p> <p>Range: 0 through 16,777,215</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring IS-IS Multicast Topology

ipv6-unicast

Syntax	ipv6-unicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure alternate IPv6 unicast topologies.</p> <p>This statement causes IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to inet6.0.</p> <div><p>NOTE: The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.</p></div>
Default	IPv6 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

ipv6-unicast-metric

Syntax	<code>ipv6-unicast-metric <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Specify the IPv6 unicast topology metric value for the level. The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics.
Options	<p><i>metric</i>—Metric value.</p> <p>Range: 0 through 16,777,215</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

isis

Syntax	isis { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable IS-IS routing on the routing device or for a routing instance. The isis statement is the one statement you must include in the configuration to run IS-IS on the routing device or in a routing instance.
Default	IS-IS is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS• Example: Configuring Multi-Level IS-IS

keep

Syntax	keep (all none);
Hierarchy Level	<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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Specify whether routes learned from a BGP peer are retained in the routing table even if they contain an AS number that was exported from the local AS.
Default	If you do not include this statement, most routes are retained in the routing table.
Options	<p>all—Retain all routes.</p> <p>none—Retain none of the routes. When keep none is configured for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Understanding Route Advertisement out-delay on page 292

labeled-unicast (Protocols BGP)

Syntax	<pre> labeled-unicast { accepted-prefix-limit { maximum <i>number</i>; teardown <<i>percentage</i>> <idle-timeout (forever <i>minutes</i>)>; } aggregate-label { community <i>community-name</i>; } explicit-null { connected-only; } prefix-limit { maximum <i>number</i>; teardown <<i>percentage</i>> <idle-timeout (forever <i>minutes</i>)>; } resolve-vpn; rib (inet.3 inet6.3); rib-group <i>group-name</i>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp <i>family</i> (inet inet6)], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <i>family</i> (inet inet6)], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> (inet inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>family</i> (inet inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> (inet inet6)], [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> <i>family</i> (inet inet6)], [edit protocols bgp <i>family</i> (inet inet6)], [edit protocols bgp group <i>group-name</i> <i>family</i> (inet inet6)], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> (inet inet6)], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>family</i> (inet inet6)], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> (inet inet6)], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> (inet inet6)] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the family type to be labeled-unicast.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

Related Documentation

- [Examples: Configuring Multiprotocol BGP](#)

level (Global IS-IS)

Syntax	<pre>level <i>level-number</i> { authentication-key <i>key</i>; authentication-key-chain (Protocols IS-IS) <i>key-chain-name</i>; authentication-type <i>type</i>; external-preference <i>preference</i>; no-csnp-authentication; no-hello-authentication; no-psnp-authentication; preference <i>preference</i>; wide-metrics-only; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>], [edit protocols <i>isis</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure the global-level properties.</p> <p>You can administratively divide a single AS into smaller groups called areas. You configure each routing device interface to be in an area. Any interface can be in any area. The area address applies to the entire routing device. You cannot specify one interface to be in one area and another interface in a different area. To route between areas, you must have two adjacent Level 2 routers that communicate with each other.</p> <p>Level 1 routers can only route within their IS-IS area. To send traffic outside their area, Level 1 routers must send packets to the nearest intra-area Level 2 router. A routing device can be a Level 1 router, a Level 2 router, or both. You specify the router level on a per-interface basis, and a routing device becomes adjacent to other routing devices on the same level on that link only.</p> <p>You can configure one Level 1 routing process and one Level 2 routing process on each interface, and you can configure the two levels differently.</p>
Options	<p><i>level-number</i>—IS-IS level number.</p> <p>Values: 1 or 2</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS• Example: Configuring Multi-Level IS-IS

link-protection (Protocols IS-IS)

Syntax	link-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable link protection on the specified IS-IS interface. Junos OS creates a backup loop-free alternate path to the primary next hop for all destination routes that traverse the protected interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Link and Node Protection for IS-IS Routes node-link-protection on page 287

local-address (Protocols BGP)

Syntax	<code>local-address address;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the address of the local end of a BGP session. This address is used to accept incoming connections to the peer and to establish connections to the remote peer. When none of the operational interfaces are configured with the specified local address, a session with a BGP peer is placed in the idle state.</p> <p>You generally configure a local address to explicitly configure the system's IP address from BGP's point of view. This IP address can be either an IPv6 or IPv4 address. Typically, an IP address is assigned to a loopback interface, and that IP address is configured here.</p> <p>For internal BGP (IBGP) peering sessions, generally the loopback interface (lo0) is used to establish connections between the IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address, the IBGP peering session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peering session also goes up and down. Thus, the loopback interface provides fault tolerance in case the physical interface or the link goes down, if the device has link redundancy.</p> <p>When a device peers with a remote device's loopback interface address, the local device expects BGP update messages to come from (be sourced by) the remote device's loopback interface address. The local-address statement enables you to specify the source information in BGP update messages. If you omit the local-address statement, the expected source of BGP update messages is based on the device's source address selection rules, which normally result in the egress interface address being the expected source of update messages. When this happens, the peering session is not established because a mismatch exists between the expected source address (the egress interface</p>

of the peer) and the actual source (the loopback interface of the peer). To ensure that the expected source address matches the actual source address, specify the loopback interface address in the **local-address** statement.



NOTE: Although a BGP session can be established when only one of the paired routers has **local-address** configured, we strongly recommend that you configure **local-address** on both paired routers for IBGP and multihop EBGP sessions. The **local-address** statement ensures that deterministic fixed addresses are used for the BGP session end-points.

If you include the **default-address-selection** statement in the configuration, the software chooses the system default address as the source for most locally generated IP packets. For protocols in which the local address is unconstrained by the protocol specification, for example IBGP and multihop EBGP, if you do not configure a specific local address when configuring the protocol, the local address is chosen using the same methods as other locally generated IP packets.

Default If you do not configure a local address, BGP uses the routing device's source address selection rules to set the local address.

Options **address**—IPv6 or IPv4 address of the local end of the connection.

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

Related Documentation

- Example: Configuring Internal BGP Peering Sessions on Logical Systems
- Example: Configuring Internal BGP Peer Sessions
- Understanding Internal BGP Peering Sessions
- [router-id on page 361](#)

local-address (Routing Options)

Syntax	<code>local-address <i>address</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast backup-pe-group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast backup-pe-group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast backup-pe-group <i>group-name</i>], [edit routing-options multicast backup-pe-group <i>group-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the address of the local PE for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.
Options	<i>address</i> —Address of local PEs in the backup group.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring Ingress PE Redundancy

local-as

Syntax	<code>local-as <i>autonomous-system</i> <loops <i>number</i>> <private alias> <no-prepend-global-as>;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>alias option introduced in Junos OS Release 9.5.</p> <p>no-prepend-global-as option introduced in Junos OS Release 9.6.</p>
Description	<p>Specify the local autonomous system (AS) number. An AS is a set of routing devices that are under a single technical administration and generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices.</p> <p>Internet service providers (ISPs) sometimes acquire networks that belong to a different AS. When this occurs, there is no seamless method for moving the BGP peers of the acquired network to the AS of the acquiring ISP. The process of configuring the BGP peers with the new AS number can be time-consuming and cumbersome. In this case, it might not be desirable to modify peer arrangements or configuration. During this kind of transition period, it can be useful to configure BGP-enabled devices in the new AS to use the former AS number in BGP updates. This former AS number is called a <i>local</i> AS.</p>



NOTE: If you are using BGP on the routing device, you must configure an AS number before you specify the local as number.

In Junos OS Release 9.1 and later, the AS numeric range in plain-number format is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For

example, the 4-byte AS number of 65546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options **alias**—(Optional) Configure the local AS as an alias of the global AS number configured for the router at the **[edit routing-options]** hierarchy level. As a result, a BGP peer considers any local AS to which it is assigned as equivalent to the primary AS number configured for the routing device. When you use the **alias** option, only the AS (global or local) used to establish the BGP session is prepended in the AS path sent to the BGP neighbor.

autonomous-system—AS number.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format

Range: 0.0 through 65535.65535 in AS-dot notation format

loops number—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.



NOTE: If you configure the local AS values for any BGP group, the detection of routing loops is performed using both the AS and the local AS values for all BGP groups.

If the local AS for the EBGP or IBGP peer is the same as the current AS, do not use the **local-as** statement to specify the local AS number.

When you configure the local AS within a VRF, this impacts the AS path loop-detection mechanism. All of the **local-as** statements configured on the device are part of a single AS domain. The AS path loop-detection mechanism is based on looking for a matching AS present in the domain.

Range: 1 through 10

Default: 1

no-prepend-global-as—(Optional) Specify to strip the global AS and to prepend only the local AS in AS paths sent to external peers.

private—(Optional) Configure to use the local AS only during the establishment of the BGP session with a BGP neighbor but to hide it in the AS path sent to external BGP peers. Only the global AS is included in the AS path sent to external peers.



NOTE: The **private** and **alias** options are mutually exclusive. You cannot configure both options with the same **local-as** statement.

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

Related Documentation

- Examples: Configuring BGP Local AS
- Example: Configuring a Local AS for EBGp Sessions
- [autonomous-system on page 94](#)
- [family on page 154](#)

local-interface (IPv6)

Syntax `local-interface interface-name;`

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp [group group-name neighbor ipv6-link-local-address](#)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* [neighbor ipv6-link-local-address](#)],
[edit protocols bgp group *group-name* [neighbor ipv6-link-local-address](#)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* [neighbor ipv6-link-local-address](#)]

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

Description Specify the interface name of the EBGp peer that uses IPv6 link-local addresses. This peer is link-local in scope.

Options *interface-name*—Interface name of the EBGp IPv6 peer.

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

Related Documentation

- Example: Configuring Internal BGP Peering Sessions on Logical Systems
- Example: Configuring Internal BGP Peer Sessions
- Example: Configuring External BGP on Logical Systems with IPv6 Interfaces
- Understanding Internal BGP Peering Sessions

local-preference

Syntax	<code>local-preference local-preference;</code>
Hierarchy Level	<code>[edit logical-systems logical-system-name protocols bgp],</code> <code>[edit logical-systems logical-system-name protocols bgp group group-name],</code> <code>[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name neighbor address],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp group group-name],</code> <code>[edit protocols bgp group group-name neighbor address],</code> <code>[edit routing-instances routing-instance-name protocols bgp],</code> <code>[edit routing-instances routing-instance-name protocols bgp group group-name],</code> <code>[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Modify the value of the LOCAL_PREF path attribute, which is a metric used by IBGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.</p> <p>The LOCAL_PREF path attribute always is advertised to internal BGP peers and to neighboring confederations. It is never advertised to external BGP peers.</p>
Default	If you omit this statement, the LOCAL_PREF path attribute, if present, is not modified.
Options	<p>local-preference—Preference to assign to routes learned from BGP or from the group or peer.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: If the LOCAL_PREF path attribute is present, do not modify its value. If a BGP route is received without a LOCAL_PREF attribute, the route is handled locally (it is stored in the routing table and advertised by BGP) as if it were received with a LOCAL_PREF value of 100. By default, non-BGP routes that are advertised by BGP are advertised with a LOCAL_PREF value of 100.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring the Local Preference Value for BGP Routes• Understanding Internal BGP Peering Sessions

- [preference on page 313](#)

log-updown (Protocols BGP)

Syntax	log-updown;
Hierarchy Level	<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 <i>group</i> <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 <i>group</i> <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 <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify to generate a log a message whenever a BGP peer makes a state transition. Messages are logged using the system logging mechanism located at the [edit system syslog] hierarchy level.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Preventing BGP Session Resets • Junos OS System Basics Configuration Guide • traceoptions on page 391

loose-authentication-check

Syntax	loose-authentication-check;
Hierarchy Level	[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]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Allow the use of MD5 authentication without requiring network-wide deployment.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Hitless Authentication Key Rollover for IS-IS

lsp-interval

Syntax	<code>lsp-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit protocols isis interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure the link-state PDU interval time.</p> <p>By default, the routing device sends one link-state PDU packet out an interface every 100 milliseconds. To disable the transmission of all link-state PDUs, set the interval to 0.</p> <p>Link-state PDU throttling by use of the lsp-interval statement controls the flooding pace to neighboring routing devices in order to not overload them.</p> <p>Also, consider that control traffic (such as link-state PDUs and related packets) might delay user traffic (information packets) because control traffic always has precedence in terms of scheduling on the routing device interface cards. Unfortunately, the control traffic transmission rate is not decreased on low-bandwidth interfaces, such as DS-0 or fractional T1 and E1 interface. Line control traffic stays the same. On a low-bandwidth circuit that is transmitting 30 full-MTU-sized packets, there is not much bandwidth left over for other types of packets.</p>
Default	By default, the routing device sends one link-state PDU out an interface every 100 milliseconds.
Options	<p><i>milliseconds</i>—Number of milliseconds between the sending of link-state PDUs. Specifying a value of 0 blocks all link-state PDU transmission.</p> <p>Range: 0 through 1000 milliseconds</p> <p>Default: 100 milliseconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces

lsp-lifetime

Syntax	<code>lsp-lifetime <i>seconds</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> <i>isis</i>],</code> <code>[edit protocols <i>isis</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Specify how long a link-state PDU originating from the routing device should persist in the network. The routing device sends link-state PDUs often enough so that the link-state PDU lifetime never expires.</p> <p>Because link-state PDUs have a maximum lifetime, they need to be refreshed. Refreshing means that a router needs to re-originate its link-state PDUs periodically. The re-origination interval must be less than the link-state PDU's lifetime. For example, if the link-state PDU is valid for 1200 seconds, the routing device needs to refresh the link-state PDU in less than 1200 seconds to avoid removal of the link-state PDU from the link-state database by other routing devices. The recommended maximum link-state PDU origination interval is the lifetime minus 300 seconds. So, in a default environment this would be 900 seconds. In Junos OS, the refresh interval is derived from the lifetime and is equal to the lifetime minus 317 seconds. You can change the lifetime to a higher value to reduce the number of refreshes in the network. (You would rarely want to increase the number of refreshes.) Often these periodic link-state PDU refreshes are referred to as refresh noise, and network administrators want to reduce this noise as much as possible.</p> <p>The <code>show isis overview</code> command displays the link-state PDU lifetime.</p>
Default	By default, link-state PDUs are maintained in network databases for 1200 seconds (20 minutes) before being considered invalid. This length of time, called the <i>LSP lifetime</i> , normally is sufficient to guarantee that link-state PDUs never expire.
Options	<i>seconds</i> —link-state PDU lifetime, in seconds. Range: 350 through 65,535 seconds Default: 1200 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces• http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

lsp-metric-into-summary

Syntax	<code>lsp-metric-into-summary;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) traffic-engineering shortcuts],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) traffic-engineering shortcuts],</p> <p>[edit protocols (ospf ospf3) traffic-engineering shortcuts],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) traffic-engineering shortcuts]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4.</p> <p>Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4 for EX Series switches.</p>
Description	Advertise the LSP metric in summary LSAs.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Support for Traffic Engineering • Example: Enabling OSPF Traffic Engineering Support

martians

Syntax	<pre>martians { destination-prefix match-type <allow>; }</pre>
Hierarchy Level	<pre>[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options], [edit logical-systems logical-system-name routing-instances routing-instance-name routing-options rib routing-table-name], [edit logical-systems logical-system-name routing-options], [edit logical-systems logical-system-name routing-options rib routing-table-name], [edit routing-instances routing-instance-name routing-options], [edit routing-instances routing-instance-name routing-options rib routing-table-name], [edit routing-options], [edit routing-options rib routing-table-name]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure martian addresses.
Options	<p>allow—(Optional) Explicitly allow a subset of a range of addresses that has been disallowed. The allow option is the only supported action.</p> <p>destination-prefix—Destination route you are configuring:</p> <ul style="list-style-type: none">• destination-prefix/prefix-length—destination-prefix is the network portion of the IP address, and prefix-length is the destination prefix length.• default—Default route to use when routing packets do not match a network or host in the routing table. This is equivalent to specifying the IP address 0.0.0.0/0. <p>match-type—Criteria that the destination must match:</p> <ul style="list-style-type: none">• exact—Exactly match the route's mask length.• longer—The route's mask length is greater than the specified mask length.• orlonger—The route's mask length is equal to or greater than the specified mask length.• through destination-prefix—The route matches the first prefix, the route matches the second prefix for the number of bits in the route, and the number of bits in the route is less than or equal to the number of bits in the second prefix.• upto prefix-length—The route's mask length falls between the two destination prefix lengths, inclusive.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- Example: Configuring Martian Addresses

max-areas

Syntax	<code>max-areas <i>number</i>;</code>
Hierarchy Level	<p>[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]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Modify the maximum number of IS-IS areas advertised.</p> <p>This value is included in the Maximum Address Area field of the IS-IS common PDU header included in all outgoing PDUs.</p> <p>The maximum number of areas you can advertise is restricted to 36 to ensure that the IIH PDUs have enough space to include other type, length, and value (TLV) fields, such as the Authentication and IPv4 and IPv6 Interface Address TLVs.</p>
Options	<p><i>number</i>—Maximum number of areas to include in the IS-IS hello (IIH) PDUs and link-state PDUs.</p> <p>Range: 3 through 36</p> <p>Default: 3</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Multi-Level IS-IS

maximum-bandwidth (Routing Options)

Syntax	maximum-bandwidth <i>bps</i> ;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> routing-instances <i>instance-name</i> routing-options multicast interface <i>interface-name</i>], [edit dynamic-profiles <i>profile-name</i> routing-options multicast interface <i>interface-name</i>] [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], [edit routing-options multicast interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. dynamic-profiles hierarchy level added in Junos OS Release 11.2.
Description	Configure the multicast bandwidth for the interface.
Options	<i>bps</i> —Bandwidth rate, in bits per second, for the multicast interface. Range: 0 through any amount of bandwidth
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Defining Interface Bandwidth Maximums

maximum-paths

Syntax	<code>maximum-paths <i>path-limit</i> <log-interval <i>seconds</i>> <log-only threshold <i>value</i>>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure a limit for the number of routes installed in a routing table based upon the route path.



NOTE: The `maximum-paths` statement is similar to the `maximum-prefixes` statement. The `maximum-prefixes` statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The `maximum-paths` limit applies the total number of routes (two). The `maximum-prefixes` limit applies to the total number of unique prefixes (one).


Options	<p><code>log-interval <i>seconds</i></code>—(Optional) Minimum time interval (in seconds) between log messages.</p> <p>Range: 5 through 86,400</p> <p><code>log-only</code>—(Optional) Sets the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.</p> <p><code><i>path-limit</i></code>—Maximum number of routes. If this limit is reached, a warning is triggered and additional routes are rejected.</p> <p>Range: 1 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: No default</p> <p><code>threshold <i>value</i></code>—(Optional) Percentage of the maximum number of routes that starts triggering a warning. You can configure a percentage of the <code><i>path-limit</i></code> value that starts triggering the warnings.</p> <p>Range: 1 through 100</p>
----------------	---



NOTE: When the number of routes reaches the threshold value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *path-limit* value, then additional routes are rejected.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs

maximum-prefixes

Syntax	<code>maximum-prefixes <i>prefix-limit</i> <log-interval <i>seconds</i>> <log-only threshold <i>percentage</i>>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Configure a limit for the number of routes installed in a routing table based upon the route prefix.</p> <p>Using a prefix limit, you can curtail the number of prefixes received from a CE router in a VPN. Prefix limits apply only to dynamic routing protocols and are not applicable to static or interface routes.</p>
<div>  <p>NOTE: The <code>maximum-prefixes</code> statement is similar to the <code>maximum-paths</code> statement. The <code>maximum-prefixes</code> statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:</p> <pre> OSPF 10.10.10.0/24 ISIS 10.10.10.0/24 </pre> <p>These are two routes, but only one destination (prefix). The <code>maximum-paths</code> limit applies the total number of routes (two). The <code>maximum-prefixes</code> limit applies to the total number of unique prefixes (one).</p> </div>	
Options	<p>log-interval <i>seconds</i>—(Optional) Minimum time interval (in seconds) between log messages.</p> <p>Range: 5 through 86,400</p> <p>log-only—(Optional) Sets the prefix limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.</p> <p><i>prefix-limit</i>—Maximum number of route prefixes. If this limit is reached, a warning is triggered and any additional routes are rejected.</p> <p>Range: 1 through 4,294,967,295</p> <p>Default: No default</p> <p>threshold <i>value</i>—(Optional) Percentage of the maximum number of prefixes that starts triggering a warning. You can configure a percentage of the <i>prefix-limit</i> value that starts triggering the warnings.</p>

Range: 1 through 100



NOTE: When the number of routes reaches the threshold value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *prefix-limit* value, then additional routes are rejected.

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

Related Documentation

- Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs

med-igp-update-interval

Syntax med-igp-update-interval *minutes*;

Hierarchy Level [edit routing-options]

Release Information Statement introduced in Junos OS Release 9.0
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure a timer for how long to delay updates for the multiple exit discriminator (MED) path attribute for BGP groups and peers configured with the **metric-out igp offset delay-med-update** statement. The timer delays MED updates for the interval configured unless the MED is lower than the previously advertised attribute or another attribute associated with the route has changed or if the BGP peer is responding to a refresh route request.

Options *minutes*—Interval to delay MED updates.
Range: 10 through 600
Default: 10 minutes

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


Related Documentation

- Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates
- [metric-out on page 253](#)

mesh-group (Protocols IS-IS)

Syntax	mesh-group (blocked <i>value</i>);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure an interface to be part of a mesh group, which is a set of fully connected nodes.</p> <p>A <i>mesh group</i> is a set of routing devices that are fully connected. That is, they have a fully meshed topology. When link-state PDUs are being flooded throughout an area, each router within a mesh group receives only a single copy of a link-state PDU instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of link-state PDUs.</p> <p>To create a mesh group and designate that an interface be part of the group, assign a mesh-group number to all the routing device interfaces in the group. To prevent an interface in the mesh group from flooding link-state PDUs, configure blocking on that interface.</p>
Options	<p>blocked—Configure the interface so that it does not flood link-state PDUs.</p> <p>value—Number that identifies the mesh group.</p> <p>Range: 1 through 4,294,967,295 ($2^{32} - 1$; 32 bits are allocated to identify a mesh group)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Mesh Groups of IS-IS Interfaces

message-size

Syntax	<code>message-size <i>number</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols <i>rip</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <i>neighbor</i></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> <i>rip</i>],</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> <i>neighbor</i> <i>neighbor-name</i>],</code> <code>[edit protocols <i>rip</i>],</code> <code>[edit protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <i>rip</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor</i></code> <code> <i>neighbor-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement for SRX Series devices introduced in Junos OS Release 9.5. Statement for J Series platform introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the number of route entries to be included in every RIP update message.
	<div><p>TIP: To ensure interoperability with other vendors' equipment, use the standard of 25 route entries per message. Do not change the default number of route entries in a RIP update message.</p></div>
Options	<i>number</i> —Number of route entries per update message. Range: 25 through 255 entries Default: 25 entries
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring RIP

metric (Protocols IS-IS)

Syntax	<code>metric <i>metric</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the metric value for the level.

All IS-IS routes have a cost, which is a routing metric that is used in the IS-IS link-state calculation. The cost is an arbitrary, dimensionless integer that can be from 1 through 63, or from 1 through 16,777,215 ($2^{24} - 1$) if you are using wide metrics.

Similar to other routing protocols, IS-IS provides a way of exporting routes from the routing table into the IS-IS network. When a route is exported into the IS-IS network without a specified metric, IS-IS uses default metric values for the route, depending on the protocol that was used to learn the route.

Table 22 on page 247 depicts IS-IS route export metric default values.

Table 22: Default Metric Values for Routes Exported into IS-IS

Protocol Used for Learning the Route	Default Metric Value
Direct	10
Static	Same as reported by the protocol used for exporting the route
Aggregate	10
Generate	10
RIP	Same as reported by the protocol used for exporting the route
OSPF	Same as reported by the protocol used for exporting the route
BGP	10

The default metric values behavior can be customized by using routing policies.

Options	<i>metric</i> —Metric value. Range: 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)
----------------	--

Default: 10 (for all interfaces except lo0), 0 (for the lo0 interface)

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

Related Documentation	<ul style="list-style-type: none">• Example: Enabling Wide IS-IS Metrics for Traffic Engineering• te-metric• wide-metrics-only on page 418
------------------------------	--

metric (Protocols OSPF Interface)

Syntax	<code>metric <i>metric</i>;</code>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> sham-link-remote], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> sham-link-remote], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Specify the cost of an OSPF interface. The cost is a routing metric that is used in the link-state calculation.</p> <p>To set the cost of routes exported into OSPF, configure the appropriate routing policy.</p>
Options	<p>metric—Cost of the route.</p> <p>Range: 1 through 65,535</p> <p>Default: By default, the cost of an OSPF route is calculated by dividing the reference-bandwidth value by the bandwidth of the physical interface. Any specific value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the cost of the route for that interface.</p>

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Controlling the Cost of Individual OSPF Network Segments• Example: Configuring OSPFv2 Sham Links• Example: Configuring Multitopology Routing Based on Applications• Example: Configuring Multitopology Routing Based on a Multicast Source• bandwidth-based-metrics on page 99• reference-bandwidth on page 334

metric (Aggregate, Generated, or Static Route)

Syntax	(metric metric2 metric3 metric4) <i>metric</i> <type type>;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)], [edit routing-options (aggregate generate static) (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specify the metric value for an aggregate, generated, or static route. You can specify up to four metric values, starting with metric (for the first metric value) and continuing with metric2 , metric3 , and metric4 .
Options	<p>metric—Metric value. Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>type type—(Optional) Type of route.</p> <p>When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric.</p> <p>Range: 1 through 16</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Summarizing Static Routes Through Route Aggregation• Example: Conditionally Generating Static Routes• aggregate on page 74• generate on page 164• static on page 376

metric-in (Protocols RIP)

Syntax	<code>metric-in <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols <code>rip</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <code>neighbor</code> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>rip</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <code>neighbor</code> <i>neighbor-name</i>],</p> <p>[edit protocols <code>rip</code>],</p> <p>[edit protocols rip group <i>group-name</i> <code>neighbor</code> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>rip</code>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <code>neighbor</code> <i>neighbor-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Specify the metric to add to incoming routes when the routing device advertises into RIP routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIP routes learned through a specific neighbor.
Options	<p><i>metric</i>—Metric value.</p> <p>Range: 1 through 16</p> <p>Default: 1</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring the Metric Value Added to Imported RIP Routes

metric-in (Protocols RIPng)

Syntax	<code>metric-in <i>metric</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols ripng],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ripng 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> ripng],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ripng group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit protocols ripng],</code> <code>[edit protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ripng],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> neighbor</code> <code> <i>neighbor-name</i>]</code>
Release Information	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.
Description	Specify the metric to add to incoming routes when advertising into RIPng routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIPng routes learned through a specific neighbor.
Options	<i>metric</i> —Metric value. Range: 1 through 16 Default: 1
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring the Metric Value Added to Imported RIPng Routes

metric-out (Protocols BGP)

Syntax	<code>metric-out (<i>metric</i> minimum-igp <i>offset</i> igp (delay-med-update <i>offset</i>);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor 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 <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name</i> <i>neighbor 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 <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Option delay-med-update introduced in Junos OS Release 9.0.</p>
Description	<p>Specify the metric for all routes sent using the multiple exit discriminator (MED, or MULTI_EXIT_DISC) path attribute in update messages. This path attribute is used to discriminate among multiple exit points to a neighboring AS. If all other factors are equal, the exit point with the lowest metric is preferred.</p> <p>You can specify a constant metric value by including the metric option. For configurations in which a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the multihop command—you can specify a variable metric by including the minimum-igp or igp option.</p> <p>You can increase or decrease the variable metric calculated from the IGP metric (either from the igp or minimum-igp statement) by specifying a value for offset. The metric is increased by specifying a positive value for offset, and decreased by specifying a negative value for offset.</p> <p>In Junos OS Release 9.0 and later, you can specify that a BGP group or peer not advertise updates for the MED path attributes used to calculate IGP costs for BGP next hops unless the MED is lower. You can also configure an interval to delay when MED updates are sent by including the med-igp-update-interval minutes statement at the [edit routing-options] hierarchy level.</p>
Options	<p>delay-med-update—Specify that a BGP group or peer configured with the metric-out igp statement not advertise MED updates unless the current MED value is lower than</p>

the previously advertised MED value, or another attribute associated with the route has changed, or the BGP peer is responding to a refresh route request.



NOTE: You cannot configure the `delay-med-update` statement at the global BGP level.

igp—Set the metric to the most recent metric value calculated in the IGP to get to the BGP next hop. Routes learned from an EBGP peer usually have a next hop on a directly connected interface and thus the IGP value is equal to zero. This is the value advertised.

metric—Primary metric on all routes sent to peers.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: No metric is sent.

minimum-igp—Set the metric to the minimum metric value calculated in the IGP to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value. When you change a neighbor's export policy from any configuration to a configuration that sets the minimum IGP offset on an exported route, the advertised MED is not updated if the value would increase as a result, even if the previous configuration does not use a minimum IGP-based MED value. This behavior helps to prevent unnecessary route flapping when an IGP cost changes, by not forcing a route update if the metric value increases past the previous lowest known value.

offset—Increases or decreases the metric by this value.

Range: -2^{31} through $2^{31} - 1$

Default: None

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

Related Documentation	<ul style="list-style-type: none">• Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates• Examples: Configuring BGP MED• Example: Configuring the MED Attribute Directly• Understanding the MED Attribute• med-igp-update-interval on page 244
------------------------------	--

metric-out (Protocols RIP)

Syntax	<code>metric-out <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rip 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 rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIP routes sent from this neighbor.</p> <p>If you have included the export statement, RIP exports routes it has learned to the neighbors configured by including the neighbor statement.</p> <p>The metric associated with a RIP route (unless modified by an export policy) is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with a metric-in value of 2 is advertised with a combined metric of 7 when advertised to RIP neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured for that group with the metric-out statement.</p> <p>The metric for a route can be modified with an export policy. That metric is seen when the route is exported to the next hop.</p> <p>To increase the metric for routes advertised outside a group, include the metric-out statement.</p>
Options	<p><i>metric</i>—Metric value.</p> <p>Range: 1 through 16</p> <p>Default: 1</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Examples: Controlling Traffic with Metrics in a RIP Network50

metric-out (Protocols RIPng)

Syntax	<code>metric-out <i>metric</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[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>],</code> <code>[edit protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>]</code>
Release Information	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.
Description	<p>Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIPng routes sent from this neighbor.</p> <p>When an export policy is configured, RIPng exports all learned routes to neighbors configured with the neighbor statement.</p> <p>If a route being exported was learned from a member of the same RIPng group, the metric associated with that route (unless modified by an export policy) is the normal RIPng metric. For example, a RIPng route with a metric of 5 learned from a neighbor configured with a metric-in value of 2 is advertised with a combined metric of 7 when advertised to RIPng neighbors in the same group. However, if this route was learned from a RIPng neighbor in a different group or from a different protocol, the route is advertised with the metric value configured for that group with the metric-out statement. The default value for metric-out is 1.</p> <p>To modify the metric for routes advertised outside a group, include the metric-out statement.</p>
Options	<i>metric</i> —Metric value. Range: 1 through 16 Default: 1
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring the Metric Value Added to Imported RIPng Routes

metric-type

Syntax	<code>metric-type type;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> nssadefault-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</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)) area <i>area-id</i> nssa default-lsa],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instances</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> nssa default-lsa]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Specify the external metric type for the default LSA.</p> <p>The configured metric determines the method used to compute the cost to a destination:</p> <ul style="list-style-type: none"> • The Type 1 external metric is equivalent to the link-state metric. The path cost uses the advertised external path cost and the path cost to the AS boundary router (the route is equal to the sum of all internal costs and the external cost). • The Type 2 external metric uses the cost assigned by the AS boundary router (the route is equal to the external cost alone). By default, OSPF uses the Type 2 external metric.
Options	type —Metric type: 1 or 2
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Areas and Router Functionality Overview • Example: Configuring OSPF Not-So-Stubby Areas

msdp

```
Syntax  msdp {
        disable;
        active-source-limit {
            log-interval seconds;
            log-warning value;
            maximum number;
            threshold number;
        }
        data-encapsulation (disable | enable);
        export [ policy-names ];
        group group-name {
            ...group-configuration ...
        }
        hold-time seconds;
        import [ policy-names ];
        local-address address;
        keep-alive seconds;
        peer address {
            ...peer-configuration ...
        }
        rib-group group-name;
        source ip-prefix</prefix-length> {
            active-source-limit {
                maximum number;
                threshold number;
            }
        }
        sa-hold-time seconds;
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
        group group-name {
            disable;
            export [ policy-names ];
            import [ policy-names ];
            local-address address;
            mode (mesh-group | standard);
            peer address {
                ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
                just following ...
            }
            traceoptions {
                file filename <files number> <size size> <world-readable | no-world-readable>;
                flag flag <flag-modifier> <disable>;
            }
        }
        peer address {
            disable;
            active-source-limit {
                maximum number;
                threshold number;
            }
        }
    }
```



```

    }
    authentication-key peer-key;
    default-peer;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
}

```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable MSDP on the router or switch. You must also configure at least one peer for MSDP to function.
Default	MSDP is disabled on the router or switch.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Configuring MSDP in a Routing Instance

mtu-discovery

Syntax	mtu-discovery;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [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>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure TCP path maximum transmission unit (MTU) discovery.</p> <p>TCP path MTU discovery enables BGP to automatically discover the best TCP path MTU for each BGP session. In Junos OS, TCP path MTU discovery is disabled by default for all BGP neighbor sessions.</p> <p>When MTU discovery is disabled, TCP sessions that are not directly connected transmit packets of 512-byte maximum segment size (MSS). These small packets minimize the chances of packet fragmentation at a device along the path to the destination. However, because most links use an MTU of at least 1500 bytes, 512-byte packets do not result in the most efficient use of link bandwidth. For directly connected EBGP sessions, MTU mismatches prevent the BGP session from being established. As a workaround, enable path MTU discovery within the EBGP group.</p> <p>Path MTU discovery dynamically determines the MTU size on the network path between the source and the destination, with the goal of avoiding IP fragmentation. Path MTU discovery works by setting the Don't Fragment (DF) bit in the IP headers of outgoing packets. When a device along the path has an MTU that is smaller than the packet, the device drops the packet. The device also sends back an ICMP Fragmentation Needed (Type 3, Code 4) message that contains the device's MTU, thus allowing the source to reduce its path MTU appropriately. The process repeats until the MTU is small enough to traverse the entire path without fragmentation.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

**Related
Documentation**

- [Example: Limiting TCP Segment Size for BGP](#)
- [Configuring the Junos OS for IPv6 Path MTU Discovery](#)
- [Configuring the Junos OS for Path MTU Discovery on Outgoing GRE Tunnel Connections](#)

multicast (Routing Options)

```
Syntax  multicast {
        forwarding-cache {
            threshold suppress value <reuse value>;
        }
        interface interface-name {
            enable;
        }
        scope scope-name {
            interface [ interface-names ];
            prefix destination-prefix;
        }
        ssm-groups {
            address;
        }
    }
```

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure generic multicast properties.



NOTE: You cannot apply a scoping policy to a specific routing instance. All scoping policies are applied to all routing instances. However, you can apply the `scope` statement to a specific routing instance.


The remaining statements are explained separately.

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

Related Documentation

- Examples: Configuring Administrative Scoping
- Example: Configuring Source-Specific Multicast Groups with Any-Source Override
- Examples: Configuring the Multicast Forwarding Cache
- Multicast Protocols Configuration Guide
- ([indirect-next-hop on page 201](#) | no-indirect-next-hop)

multihop

Syntax	<pre>multihop { no-nexthop-change; ttl <i>ttl-value</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [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>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure an EBGp multihop session.</p> <p>For Layer 3 VPNs, you configure the EBGp multihop session between the PE and CE routers. This allows you to configure one or more routers between the PE and CE routers.</p> <p>An external confederation peer is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case because multihop behavior is implied.</p> <p>If you have external BGP confederation peer-to-loopback addresses, you still need the multihop configuration.</p> <div style="margin-top: 20px;">  <p>NOTE: You cannot configure the <code>accept-remote-nexthop</code> statement at the same time.</p> </div>
Default	<p>If you omit this statement, all EBGp peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or “regular,” BGP session), and the default time-to-live (TTL) value is 1.</p> <p>The remaining statements are explained separately.</p>

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring EBGp Multihop Sessions• Configuring EBGp Multihop Sessions Between PE and CE Routers in Layer 3 VPNs• accept-remote-nextthop on page 69• no-nextthop-change• ttl

multipath (Protocols BGP)

Syntax	<pre> multipath { multiple-as; vpn-unequal-cost equal-external-internal; } </pre>
Hierarchy Level	<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 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 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 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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Allow load sharing among multiple EBGP paths and multiple IBGP paths. A path is considered a BGP equal-cost path (and will be used for forwarding) if a tie-break is performed. The tie-break is performed after the BGP route path selection step that chooses the next-hop path that is resolved through the IGP route with the lowest metric. All paths with the same neighboring AS, learned by a multipath-enabled BGP neighbor, are considered.</p>
Options	<p>multiple-as—Disable the default check requiring that paths accepted by BGP multipath must have the same neighboring AS.</p> <p>vpn-unequal-cost equal-external-internal—Enable load-balancing in a Layer 3 VPN with unequal cost paths.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Understanding BGP Path Selection Example: Load Balancing BGP Traffic

neighbor (Protocols BGP)

```
Syntax  neighbor address {
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    as-override;
    authentication-algorithm algorithm;
    authentication-key key;
    authentication-key-chain key-chain;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family {
        (inet | inet6 | inet-mvpn | inet6-mpvn | inet-vpn | inet6-vpn | iso-vpn | l2-vpn) {
            (any | flow | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                damping;
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                rib-group group-name;
                topology name {
                    community {
                        target identifier;
                    }
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib inet.3;
        }
    }
}
```



```

    rib-group group-name;
    topology name {
        community {
            target identifier;
        }
    }
}
route-target {
    advertise-default;
    external-paths number;
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
signaling {
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-interface interface-name;
local-preference preference;
log-updown;
metric-out (metric | minimum-igp <offset> | igp <offset>);
mtu-discovery;
multihop <ttl-value>;
multipath {
    multiple-as;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
}

```

```
    flag flag <flag-modifier> <disable>;  
  }  
  vpn-apply-export;  
}
```

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Explicitly configure a neighbor (peer). To configure multiple BGP peers, include multiple **neighbor** statements.

By default, the peer's options are identical to those of the group. You can override these options by including peer-specific option statements within the **neighbor** statement.

The **neighbor** statement is one of the statements you can include in the configuration to define a minimal BGP configuration on the routing device. (You can include an **allow all** statement in place of a **neighbor** statement.)

Options **address**—IPv6 or IPv4 address of a single peer.

The remaining statements are explained separately.

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

Related Documentation

- BGP Configuration Guide

neighbor (Protocols RIP)

Syntax	<pre> neighbor <i>neighbor-name</i> { authentication-key <i>password</i>; authentication-type <i>type</i>; bfd-liveness-detection { authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check; } detection-time { threshold <i>milliseconds</i>; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; transmit-interval { threshold <i>milliseconds</i>; minimum-interval <i>milliseconds</i>; } multiplier <i>number</i>; version (0 1 automatic); } (<i>check-zero</i> <i>no-check-zero</i>); demand-circuit; import <i>policy-name</i>; max-retrans-time <i>seconds</i>; message-size <i>number</i>; metric-in <i>metric</i>; metric-out <i>metric</i>; receive <i>receive-options</i>; route-timeout <i>seconds</i>; send <i>send-options</i>; update-interval <i>seconds</i>; }</pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols rip <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <i>group</i> <i>group-name</i>], [edit protocols rip <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols rip <i>group</i> <i>group-name</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Configure neighbor-specific RIP parameters, thereby overriding the defaults set for the routing device.
Options	<p><i>neighbor-name</i>—Name of an interface over which a routing device communicates to its neighbors.</p> <p>The remaining statements are explained separately.</p>

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

Related Documentation

- Example: Configuring RIP

neighbor (Protocols RIPng)

Syntax `neighbor neighbor-name {
 import [policy-names];
 metric-in metric;
 receive <none>;
 route-timeout seconds;
 send <none>;
 update-interval seconds;
}`

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

Release Information 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.

Description Configure neighbor-specific RIPng parameters, thereby overriding the defaults set for the routing device.

Options *neighbor-name*—Name of an interface over which a routing device communicates to its neighbors.

The remaining statements are explained separately.

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

Related Documentation

- Example: Configuring RIPng

no-adjacency-holddown

Syntax	no-adjacency-holddown;
Hierarchy Level	[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]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Disable the hold-down timer for IS-IS adjacencies. A hold-down timer delays the advertising of adjacencies by waiting until a time period has elapsed before labeling adjacencies in the up state. You can disable this hold-down timer, which labels adjacencies up faster. However, disabling the hold-down timer creates more frequent link-state PDU updates and SPF computation.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • hold-time on page 189

no-aggregator-id

Syntax	no-aggregator-id;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [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>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Prevent different routers within an AS from creating aggregate routes that contain different AS paths.</p> <p>Junos OS performs route aggregation, which is the process of combining the characteristics of different routes so that only a single route is advertised. Aggregation reduces the amount of information that BGP must store and exchange with other BGP systems. When aggregation occurs, the local routing device adds the local AS number and the router ID to the aggregator path attribute. The no-aggregator-id statement causes Junos OS to place a 0 in the router ID field and thus eliminate the possibility of having multiple aggregate advertisements in the network, each with different path information.</p>
Default	If you omit this statement, the router ID is included in the BGP aggregator path attribute.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• BGP Messages Overview

no-authentication-check

Syntax	no-authentication-check;
Hierarchy Level	[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]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Generate authenticated packets and check the authentication on received packets, but do not reject packets that cannot be authenticated.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• hello-authentication-type on page 181

no-client-reflect

Syntax	no-client-reflect;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [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>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Disable intracluster route redistribution by the system acting as the route reflector. Include this statement when the client cluster is fully meshed to prevent the sending of redundant route advertisements. Route reflection provides a way to decrease BGP control traffic and minimizing the number of update messages sent within the AS.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BGP Route Reflectors• cluster on page 125

no-csnp-authentication

Syntax	no-csnp-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Suppress authentication check on complete sequence number PDU (CSNP) packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • csnp-interval on page 130

no-eligible-backup (Protocols IS-IS)

Syntax	no-eligible-backup;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Exclude the specified interface as a backup interface for IS-IS interfaces on which link protection or node-link protection is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Link and Node Protection for IS-IS Routes • link-protection on page 225 • node-link-protection on page 287


no-hello-authentication

Syntax	no-hello-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Suppress authentication check on complete sequence number hello packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• hello-authentication-type on page 181

no-ipv4-multicast

Syntax	no-ipv4-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Exclude an interface from IPv4 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS Multicast Topology

no-ipv4-routing

Syntax	no-ipv4-routing;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Disable IP version 4 (IPv4) routing.</p> <p>Disabling IPv4 routing has the following results:</p> <ul style="list-style-type: none"> • The routing device does not advertise the network layer protocol identifier (NLPID) for IPv4 in the Junos OS link-state PDU fragment zero. • The routing device does not advertise any IPv4 prefixes in Junos OS link-state PDUs. • The routing device does not advertise the NLPID for IPv4 in Junos OS hello packets. • The routing device does not advertise any IPv4 addresses in Junos OS hello packets. • The routing device does not calculate any IPv4 routes.
	<div>  <p>NOTE: Note: Even when <code>no-ipv4-routing</code> is configured, an IS-IS traceoptions log can list rejected IPv4 addresses. When a configuration is committed, IS-IS schedules a scan of the routing table to determine whether any routes need to be exported into the IS-IS link state database. The implicit default export policy action is to reject everything. IPv4 addresses from the routing table are examined for export, rejected by the default policy, and the rejections are logged.</p> </div>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

no-ipv6-multicast

Syntax	no-ipv6-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Exclude an interface from the IPv6 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS Multicast Topology

no-ipv6-routing

Syntax	no-ipv6-routing;
Hierarchy Level	[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]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Disable IP version 6 (IPv6) routing. Disabling IPv6 routing has the following results: <ul style="list-style-type: none"> • The routing device does not advertise the network layer protocol identifier (NLPID) for IPv6 in the Junos OS link-state PDU fragment zero. • The routing device does not advertise any IPv6 prefixes in Junos OS link-state PDUs. • The routing device does not advertise the NLPID for IPv6 in Junos OS hello packets. • The routing device does not advertise any IPv6 addresses in Junos OS hello packets. • The routing device does not calculate any IPv6 routes.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

no-ipv6-unicast

Syntax	no-ipv6-unicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Exclude an interface from the IPv6 unicast topologies. This enables you to exercise control over the paths that unicast data takes through a network.
Default	IPv6 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

no-nssa-abr

Syntax	no-nssa-abr;
Hierarchy Level	<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>
Release Information	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm 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>
Description	Disable exporting Type 7 link-state advertisements into not-so-stubby-areas (NSSAs) for an autonomous system boundary router (ASBR) or an area border router (ABR).
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring OSPF Not-So-Stubby Areas

no-psnp-authentication

Syntax	no-psnp-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Suppress authentication check on partial sequence number PDU (PSNP) packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring IS-IS Authentication

no-qos-adjust

Syntax	no-qos-adjust;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i> reverse-oif-mapping],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i> reverse-oif-mapping],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i> reverse-oif-mapping],</p> <p>[edit routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-options multicast interface <i>interface-name</i> reverse-oif-mapping]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Statement added to [edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], and [edit routing-options multicast interface <i>interface-name</i>] hierarchy levels in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Disable hierarchical bandwidth adjustment for all subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Multicast with Subscriber VLANs

no-rfc-1583

Syntax	no-rfc-1583;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (ospf ospf3)], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Disable compatibility with RFC 1583, <i>OSPF Version 2</i> . If the same external destination is advertised by AS boundary routers that belong to different OSPF areas, disabling compatibility with RFC 1583 can prevent routing loops.
Default	Compatibility with RFC 1583 is enabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Disabling OSPFv2 Compatibility with RFC 1583

no-unicast-topology

Syntax	no-unicast-topology;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Exclude an interface from the IPv4 unicast topologies.
Default	IPv4 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring IS-IS Multicast Topology

no-validate

Syntax	<code>no-validate <i>policy-name</i>;</code>
Hierarchy Level	<code>[edit protocols bgp group <i>group-name</i> family (inet inet flow)],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor address family (inet inet flow)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet flow)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor address family (inet inet flow)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>When BGP is carrying flow-specification network layer reachability information (NLRI) messages, the no-validate statement omits the flow route validation procedure after packets are accepted by a policy.</p> <p>The receiving BGP-enabled device accepts a flow route if it passes the following criteria:</p> <ul style="list-style-type: none">• The originator of a flow route matches the originator of the best match unicast route for the destination address that is embedded in the route.• There are no more specific unicast routes, when compared to the destination address of the flow route, for which the active route has been received from a different next-hop autonomous system. <p>The first criterion ensures that the filter is being advertised by the next-hop used by unicast forwarding for the destination address embedded in the flow route. For example, if a flow route is given as 10.1.1.1, proto=6, port=80, the receiving BGP-enabled device selects the more specific unicast route in the unicast routing table that matches the destination prefix 10.1.1.1/32. On a unicast routing table containing 10.1/16 and 10.1.1/24, the latter is chosen as the unicast route to compare against. Only the active unicast route entry is considered. This follows the concept that a flow route is valid if advertised by the originator of the best unicast route.</p> <p>The second criterion addresses situations in which a given address block is allocated to different entities. Flows that resolve to a best-match unicast route that is an aggregate route are only accepted if they do not cover more specific routes that are being routed to different next-hop autonomous systems.</p> <p>You can bypass the validation process and use your own specific import policy. To disable the validation procedure and use an import policy instead, include the no-validate statement in the configuration.</p> <p>Flow routes configured for VPNs with family inet-vpn are not automatically validated, so the no-validate statement is not supported at the <code>[edit protocols bgp group <i>group-name</i> family inet-vpn]</code> hierarchy level. No validation is needed if the flow routes are configured locally between devices in a single AS.</p>

Options	<i>policy-name</i> —Import policy to match NLRI messages.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Flow Routes

node-link-protection (Protocols IS-IS)

Syntax	node-link-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable node-link protection on the specified IS-IS interface. Junos OS creates an alternate loop-free path to the primary next hop for all destination routes that traverse a protected interface. This alternate path avoids the primary next-hop routing device altogether and establishes a path through a different routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Link and Node Protection for IS-IS Routes • link-protection on page 225

nssa

Syntax	<pre>nssa { area-range <i>network/mask-length</i> <restrict> <exact> <override-metric <i>metric</i>>; default-lsa { default-metric <i>metric</i>; metric-type <i>type</i>; type-7; } (no-summaries summaries); }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (ospf ospf3) area <i>area-id</i>], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Configure a not-so-stubby area (NSSA). An NSSA allows external routes to be flooded within the area. These routes are then leaked into other areas.</p> <p>You cannot configure an area as being both a stub area and an NSSA.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• OSPF Areas and Router Functionality Overview• Example: Configuring OSPF Not-So-Stubby Areas• stub on page 382

options (Routing Options)

Syntax	<pre>options { syslog (level <i>level</i> upto level <i>level</i>); }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Configure the types of system logging messages sent about the routing protocols process to the system message logging file. These messages are also displayed on the system console. You can log messages at a particular level, or up to and including a particular level.</p>
Options	<p>level <i>level</i>—Severity of the message. It can be one or more of the following levels, in order of decreasing urgency:</p> <ul style="list-style-type: none"> • alert—Conditions that should be corrected immediately, such as a corrupted system database. • critical—Critical conditions, such as hard drive errors. • debug—Software debugging messages. • emergency—Panic or other conditions that cause the system to become unusable. • error—Standard error conditions. • info—Informational messages. • notice—Conditions that are not error conditions, but might warrant special handling. • warning—System warning messages. <p>upto level <i>level</i>—Log all messages up to a particular level.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • syslog in the Junos OS System Basics Configuration Guide

ospf

Syntax	ospf { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable OSPF routing on the routing device. You must include the ospf statement to enable OSPF on the routing device.
Default	OSPF is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• OSPF Configuration Overview• [edit protocols ospf] Hierarchy Level


ospf3

Syntax	ospf3 { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable OSPFv3 routing on the routing device. You must include the ospf3 statement to enable OSPFv3.
Default	OSPFv3 is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • OSPF Configuration Overview • [edit protocols ospf3] Hierarchy Level

out-delay

Syntax	<code>out-delay seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i></code> <code> <i>neighbor</i> <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 <i>group</i> <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 <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group</i> <i>group-name</i>],</code> <code>[edit protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <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 <i>group</i> <i>group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i></code> <code> <i>neighbor</i> <i>address</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Specify how long a route must be present in the Junos OS routing table before it is exported to BGP. Use this time delay to help bundle routing updates. When configured, the out-delay value displays as Outbound Timer when using show bgp group or show bgp group neighbor commands.
Default	If you omit this statement, routes are exported to BGP immediately after they have been added to the routing table.
Options	seconds —Output delay time. Range: 0 through 65,535 seconds Default: 0 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Understanding Route Advertisement

outbound-route-filter

Syntax	<pre> outbound-route-filter { bgp-orf-cisco-mode; prefix-based { accept { (inet inet6); } } } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [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>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>] </pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>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>
Description	Configure a BGP peer to accept outbound route filters from a remote peer.
Options	<p>accept—Specify that outbound route filters from a BGP peer be accepted.</p> <p>inet—Specify that IPv4 prefix-based outbound route filters be accepted.</p> <p>inet6—Specify that IPv6 prefix-based outbound route filters be accepted.</p>
	<div>  <p>NOTE: You can specify that both IPv4 and IPv6 outbound route filters be accepted.</p> </div>
	<p>prefix-based—Specify that prefix-based filters be accepted.</p> <p>The bgp-orf-cisco-mode statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Configuring BGP Prefix-Based Outbound Route Filtering](#)

overload (Protocols IS-IS)

Syntax	<pre> overload { advertise-high-metrics; allow-route-leaking; timeout <i>seconds</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>], [edit protocols <i>isis</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure the local routing device so that it appears to be overloaded. This statement causes the routing device to continue participating in IS-IS routing, but prevents it from being used for transit traffic. Traffic destined to immediately attached subnets continues to transit the routing device.</p> <p>You can also advertise maximum link metrics in network layer reachability information (NLRI) instead of setting the overload bit.</p> <p>You configure or disable overload mode in IS-IS with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the IS-IS instance started is less than the specified timeout.</p> <p>A timer is started for the difference between the timeout and the time elapsed since the instance started. If the time elapsed after the IS-IS instance is enabled is less than the specified timeout, overload mode is set. When the timer expires, overload mode is cleared. In overload mode, the routing device IS-IS advertisements are originated with the overload bit set. This causes the transit traffic to take paths around the routing device. However, the overloaded routing device's own links are still accessible.</p> <p>The value of the overload bit depends on these three scenarios:</p> <ol style="list-style-type: none"> 1. When the overload bit has already been set to a given value and the routing process is restarted: Link-state PDUs are regenerated with the overload bit cleared. 2. When the overload bit is reset to a lesser value while the routing process is running: Link-state PDUs are regenerated with the overload bit cleared. 3. When the overload bit is reset to a greater value while the routing process is running: Link-state PDUs are regenerated with the overload bit set to the difference between the old and new value. <p>In overload mode, the routing device advertisement is originated with all the transit routing device links (except stub) set to a metric of 0xFFFF. The stub routing device links are</p>

advertised with the actual cost of the interfaces corresponding to the stub. This causes the transit traffic to avoid the overloaded routing device and take paths around the routing device.

To understand the reason for setting the overload bit, consider that BGP converges slowly. It is not very good at detecting that a neighbor is down because it has slow-paced keepalive timers. Once the BGP neighbor is determined to be down, it can take up to 2 minutes for a BGP router to declare the neighbor down. IS-IS is much quicker. IS-IS only takes 10-30 seconds to detect absent peers. It is the slowness of BGP, more precisely the slowness of internal BGP (IBGP), that necessitates the use of the overload bit. IS-IS and BGP routing are mutually dependent on each other. If both do not converge at the same time, traffic is dropped without notification (black holed).

You might want to configure the routing device so that it appears to be overloaded when you are restarting routing on the device. Setting the overload bit for a fixed amount of time right after a restart of the routing protocol process (rpd) ensures that the router does not receive transit traffic while the routing protocols (especially IBGP) are still converging.

Setting the overload bit is useful when performing hardware or software maintenance work on a routing device. After the maintenance work, clear the overload bit to carry on forwarding transit traffic. Manual clearing of the overload bit is not always possible. What is needed is an automated way of clearing the overload bit after some amount of time. Most networks use a time value of 300 seconds. This 5-minute value provides a good balance, allowing time to bring up even large internal IBGP meshes, while still relatively quick.

Another appropriate application for setting for the overload bit is on dedicated devices such as BGP route reflectors, which are intentionally not meant to carry any transit traffic. In this case, you would not use the timer.

You can verify that the overload bit is set by running the **show isis database** command.

Options **advertise-high-metrics**—Advertise maximum link metrics in NLRI instead of setting the overload bit.

When you configure the **advertise-high-metrics** option, the routing device in overload mode stops passing (leaking) route information into the network. So an L1-L2 routing device in overload mode stops passing route information between Level 1 and Level 2 and clears its attached bit when the **advertise-high-metrics** option is configured.

Default: With **advertise-high-metrics** configured, the routing device in overload mode stops leaking route information into the network.

allow-route-leaking—Enable leaking of route information into the network even if the overload bit is set.



NOTE: The **allow-route-leaking** option does not work if the routing device is in dynamic overload mode. Dynamic overload can occur if the device has exceeded its resource limits, such as the prefix limit.

timeout seconds—Number of seconds at which the overloading is reset.

Range: 60 through 1800 seconds


Default: 0 seconds

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

Related Documentation

- Example: Configuring IS-IS

overload (Protocols OSPF)

Syntax	<pre>overload { timeout <i>seconds</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)],</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 ospf topology (default ipv4-multicast <i>name</i>)],</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 ospf topology (default ipv4-multicast <i>name</i>)],</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 ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm 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>
Description	<p>Configure the local routing device so that it appears to be overloaded. You might do this when you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic.</p>
	<div>  <p>NOTE: Traffic destined to directly attached interfaces continues to reach the routing device.</p> </div>
Options	<p>timeout <i>seconds</i>—(Optional) Number of seconds at which the overloading is reset. If no timeout interval is specified, the routing device remains in overload state until the overload statement is deleted or a timeout is set.</p> <p>Range: 60 through 1800 seconds</p> <p>Default: 0 seconds</p>




NOTE: Multitopology Routing does not support the timeout option.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring OSPF to Make Routing Devices Appear Overloaded • Example: Configuring Multitopology Routing Based on Applications • Example: Configuring Multitopology Routing Based on a Multicast Source

passive (Protocols BGP)

Syntax	passive;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [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>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure the router so that active open messages are not sent to the peer. Once you configure the routing device to be passive, the routing device will wait for the peer to issue an open request before a message is sent.
Default	If you omit this statement, all explicitly configured peers are active, and each peer periodically sends open requests until its peer responds.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Preventing BGP Session Flaps When VPN Families Are Configured

passive (Protocols IS-IS)

Syntax	<code>passive;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Advertise the direct interface addresses on an interface or into a level on the interface without actually running IS-IS on that interface or level.</p> <p>This statement effectively prevents IS-IS from running on the interface. To enable IS-IS on an interface, include the interface statement at the [edit protocols isis] or the [edit routing-instances <i>routing-instance-name</i> protocols isis] hierarchy level. To disable it, include the disable statement at those hierarchy levels. The three states—enabling, disabling, or not running IS-IS on an interface—are mutually exclusive.</p>
	<p> NOTE: Configuring IS-IS on a loopback interface automatically renders it as a passive interface, irrespective of whether the passive statement was used in the configuration of the interface.</p>
	<p>If neither passive mode nor the family iso option is configured on the IS-IS interface, then the routing device treats the interface as not being operational, and no direct IPv4/IPv6 routes are exported into IS-IS. (You configure the family iso option at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.)</p>
Default	By default, IS-IS must be configured on an interface or a level for direct interface addresses to be advertised into that level.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Multi-Level IS-IS disable

passive (Protocols OSPF)

Syntax	<pre> passive { traffic-engineering { remote-node-id address; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</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) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>traffic-engineering and remote-node-id address statements introduced in Junos OS Release 8.0.</p> <p>traffic-engineering and remote-node-id address statements introduced in Junos OS Release 8.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Advertise the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run.</p> <p>To configure an interface in OSPF passive traffic engineering mode, include the traffic-engineering statement. Configuring OSPF passive traffic engineering mode enables the dynamic discovery of OSPF AS boundary routers.</p> <p>Enable OSPF on an interface by including the interface statement at the [edit protocols (ospf ospf3) area <i>area-id</i>] or the [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i>] hierarchy levels. Disable it by including the disable statement. To prevent OSPF from running on an interface, include the passive statement. These three states are mutually exclusive.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

**Related
Documentation**

- Example: Configuring a Passive OSPF Interface
- Example: Configuring OSPF Passive Traffic Engineering Mode
- [disable on page 139](#)

peer-as (Protocols BGP)

Syntax	<code>peer-as <i>autonomous-system</i>;</code>
Hierarchy Level	<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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the neighbor (peer) autonomous system (AS) number.</p> <p>For EBGP, the peer is in another AS, so the AS number you specify in the peer-as statement must be different from the local router's AS number, which you specify in the autonomous-system statement. For IBGP, the peer is in the same AS, so the two AS numbers that you specify in the autonomous-system and peer-as statements must be the same.</p> <p>The AS numeric range in plain-number format has been extended in Junos OS Release 9.1 to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of the Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p> <p>With the introduction of 4-byte AS numbers, you might have a combination of routers that support 4-byte AS numbers and 2-byte AS numbers. For more information about what happens when establishing BGP peer relationships between 4-byte and 2-byte capable routers, see the following topics:</p>

- Establishing a Peer Relationship Between a 4-Byte Capable Router and a 2-Byte Capable Router Using a 2-Byte AS Number in the *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.
- Establishing a Peer Relationship Between a 4-Byte Capable Router and a 2-Byte Capable Router Using a 4-Byte AS Number in the *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.

Options *autonomous-system*—AS number.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format for 4-byte AS numbers

Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte AS numbers

Required Privilege routing—To view this statement in the configuration.

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

- Related Documentation**
- 4-Byte Autonomous System Numbers Overview in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)
 - Juniper Networks Implementation of 4-Byte Autonomous System Numbers in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)

pim-to-igmp-proxy

Syntax	<code>pim-to-igmp-proxy { upstream-interface [interface-names]; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 9.6 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.
Description	<p>Configure the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Internet Group Management Protocol (IGMP) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the pim-to-igmp-proxy statement enables you to use IGMP to forward IPv4 multicast traffic across the PIM sparse mode domains.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring PIM-to-IGMP Message Translation

pim-to-mld-proxy

Syntax	<pre>pim-to-mld-proxy { upstream-interface [interface-names]; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Configure the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Multicast Listener Discovery (MLD) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the pim-to-mld-proxy statement enables you to use MLD to forward IPv6 multicast traffic across the PIM sparse mode domains.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring PIM-to-MLD Message Translation

point-to-point

Syntax	point-to-point;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure an IS-IS interface to behave like a point-to-point connection.</p> <p>You can use the point-to-point statement to configure a LAN interface to act like a point-to-point interface for IS-IS. You do not need an unnumbered LAN interface, and it has no effect if configured on an interface that is already point-to-point.</p> <p>The point-to-point statement affects only IS-IS protocol procedures on that interface. All other protocols continue to treat the interface as a LAN interface. Only two IS-IS routing devices can be connected to the LAN interface, and both must be configured as point-to-point.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> IS-IS Overview Understanding IS-IS Designated Routers Example: Configuring Synchronization Between IS-IS and LDP

policy (Aggregate and Generated Routes)

Syntax	<code>policy <i>policy-name</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)],</p> <p>[edit routing-options (aggregate generate) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Associate a routing policy when configuring an aggregate or generated route's destination prefix in the routes part of the aggregate or generate statement. This provides the equivalent of an import routing policy filter for the destination prefix. That is, each potential contributor to an aggregate route, along with any aggregate options, is passed through the policy filter. The policy then can accept or reject the route as a contributor to the aggregate route.</p> <p>If the contributor is accepted, the policy can modify the default preferences. The contributor with the numerically smallest prefix becomes the most preferred, or <i>primary</i>, contributor. A rejected contributor still can contribute to a less specific aggregate route. If you do not specify a policy filter, all candidate routes contribute to an aggregate route.</p> <p>The following algorithm is used to compare two generated contributing routes in order to determine which one is the primary or preferred contributor:</p> <ol style="list-style-type: none"> 1. Compare the protocol's preference of the contributing routes. The lower the preference, the better the route. This is similar to the comparison that is done while determining the best route for the routing table. 2. Compare the protocol's preference2 of the contributing routes. The lower preference2 value is better. If only one route has preference2, then this route is preferred. 3. The preference values are the same. Proceed with a numerical comparison of the prefixes' values. <ol style="list-style-type: none"> a. The primary contributor is the numerically smallest prefix value. b. If the two prefixes are numerically equal, the primary contributor is the route that has the smallest prefix length value.

At this point, the two routes are the same. The primary contributor does not change. An additional next hop is available for the existing primary contributor.

A rejected contributor still can contribute to less specific generated route. If you do not specify a policy filter, all candidate routes contribute to a generated route.

Options	<i>policy-name</i> —Name of a routing policy.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Summarizing Routes Through Route Aggregation • Example: Conditionally Generating Static Routes • aggregate on page 74 • generate on page 164

policy (Flow Maps)

Syntax	<code>policy [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-options multicast flow-map <i>flow-map-name</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure a flow map policy.
Options	<i>policy-names</i> —Name of one or more policies for flow mapping.
Required Privilege Level	routing—To view this statement in the configuration.

policy (SSM Maps)

Syntax	<code>policy [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>],</code> <code>[edit routing-options multicast ssm-map <i>ssm-map-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Apply one or more policies to an SSM map.
Options	<i>policy-names</i> —Name of one or more policies for SSM mapping.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To view this statement in the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring SSM Mapping

ppm (Ethernet Switching)

Syntax	ppm { centralized; }
Hierarchy Level	[edit protocols lacp]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.1 for T Series devices.
Description	Configure PPM processing options for Link Aggregation Control Protocol (LACP) packets. This command configures the PPM processing options for LACP packets only. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the no-delegate-processing configuration statement in the [edit routing-options ppm] statement hierarchy.
Default	Distributed PPM processing is enabled for all packets that use PPM.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure) on page 31 • Configuring Distributed Periodic Packet Management

ppm

Syntax	<pre>ppm { no-delegate-processing; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 10.2 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.
Description	<p>(M120, M320, MX Series, T Series, TX Matrix routers, M7i and M10i routers with Enhanced CFEB [CFEB-E], EX Series switches, and QFX Series only) Disable distributed periodic packet management (PPM) to the Packet Forwarding Engine (on routers), to access ports (on EX3200 and EX4200 switches, and QFX Series), or to line cards (on EX6200 and EX8200 switches).</p> <p>After you disable PPM, PPM processing continues to run on the Routing Engine.</p> <p>In Junos OS Release 8.2, PPM was moved from the Routing Engine to the Packet Forwarding Engine, access ports, or line cards. The no-delegate-processing statement disables the default behavior and restores the legacy behavior.</p>
Default	Distributed PPM processing is enabled for all protocols that use PPM.
Options	no-delegate-processing —Disable PPM to the Packet Forwarding Engine, access ports, or line cards. Distributed PPM is enabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure) on page 31• Configuring Distributed Periodic Packet Management

preference (Protocols BGP)

Syntax	<code>preference preference;</code>
Hierarchy Level	<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 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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the preference for routes learned from BGP.</p> <p>At the BGP global level, the preference statement sets the preference for routes learned from BGP. You can override this preference in a BGP group or peer preference statement.</p> <p>At the group or peer level, the preference statement sets the preference for routes learned from the group or peer. Use this statement to override the preference set in the BGP global preference statement when you want to favor routes from one group or peer over those of another.</p>
Options	<p>preference—Preference to assign to routes learned from BGP or from the group or peer.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 170 for the primary preference</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • local-preference on page 232 • Example: Configuring the Preference Value for BGP Routes

preference (Protocols IS-IS)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure the preference of internal routes.</p> <p>Route preferences (also known as administrative distances) are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.</p> <p>To change the preference values, include the preference statement (for internal routes) or the external-preference statement.</p>
Options	<p>preference—Preference value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Route Preferences Overview• Example: Redistributing OSPF Routes into IS-IS• Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS• external-preference on page 152

preference (Protocols OSPF)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	<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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm 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>
Description	Set the route preference for OSPF internal routes.
Options	<p><i>preference</i>—Preference value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 10</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Controlling OSPF Route Preferences external-preference on page 153

preference (Protocols RIP)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols rip 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> rip group <i>group-name</i>],</code> <code>[edit protocols rip group <i>group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Specify the preference of external routes learned by RIP as compared to those learned from other routing protocols.</p> <p>By default, Junos OS assigns a preference of 100 to routes that originate from RIP. When Junos OS determines a route's preference to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.</p>
Options	<p><i>preference</i>—Preference value. A lower value indicates a more preferred route.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 100</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Route Preferences Overview

preference (Protocols RIPng)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ripng group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i>],</p> <p>[edit protocols ripng group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>
Description	<p>Specify the preference of external routes learned by RIPng as compared to those learned from other routing protocols.</p> <p>By default, Junos OS assigns a preference of 100 to routes that originate from RIPng. When Junos OS determines that a route is to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.</p> <p>To modify the default RIPng preference value, include the preference statement.</p>
Options	<p>preference—Preference value. A lower value indicates a more preferred route.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 100</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring RIPng

preference (Routing Options)

Syntax	<code>(preference preference2 color color2) preference <type type>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Preference value for a static, aggregate, or generated route. You also can specify a secondary preference value (preference2), as well as colors, which are even finer-grained preference values (color and color2).</p> <p>If the Junos OS routing table contains a dynamic route to a destination that has a better (lower) preference value than the static, aggregate, or generated route, the dynamic route is chosen as the active route and is installed in the forwarding table.</p>
Options	<p>preference—Preference value. A lower number indicates a more preferred route.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 5 (for static routes), 130 (for aggregate and generated routes)</p> <p>type type—(Optional) Type of route.</p> <p>Range: 1 through 16</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Examples: Configuring Static Routes • Example: Summarizing Routes Through Route Aggregation • Example: Conditionally Generating Static Routes • aggregate on page 74 • generate on page 164 • static on page 376

prefix

Syntax	<code>prefix destination-prefix;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast scope <i>scope-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast scope <i>scope-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast scope <i>scope-name</i>],</p> <p>[edit routing-options multicast scope <i>scope-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Configure the prefix for multicast scopes.
Options	destination-prefix —Address range for the multicast scope.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Examples: Configuring Administrative Scoping Example: Creating a Named Scope for Multicast Scoping multicast

prefix-export-limit (Protocols IS-IS)

Syntax	<code>prefix-export-limit <i>number</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols isis level level-number],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>isis level level-number],</code> <code>[edit protocols isis level level-number],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure a limit to the number of prefixes exported into IS-IS.</p> <p>By default, there is no limit to the number of prefixes that can be exported into IS-IS. To configure a limit to the number of prefixes that can be exported into IS-IS, include the prefix-export-limit statement. The prefix-export-limit statement protects the rest of the network from a malicious policy by applying a threshold filter for exported routes.</p> <p>The number of prefixes depends on the size of your network. Good design advice is to set it to double the total number of IS-IS Level 1 and Level 2 routing devices in your network.</p> <p>If the number of prefixes exported into IS-IS exceeds the configured limit, the overload bit is set and the overload state is reached. When other routers detect that this bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes. The overload state can be cleared by using the clear isis overload command.</p> <p>The show isis overview command displays the prefix export limit when it is configured.</p>
Options	<p><i>number</i>—Prefix limit.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS• Example: Redistributing OSPF Routes into IS-IS

prefix-export-limit (Protocols OSPF)

Syntax	<code>prefix-export-limit <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)],</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 ospf topology (default ipv4-multicast <i>name</i>)],</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 ospf topology (default ipv4-multicast <i>name</i>)],</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 ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm 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>
Description	Configure a limit to the number of prefixes exported into OSPF.
Options	<p><i>number</i>—Prefix limit.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Limiting the Number of Prefixes Exported to OSPF • Example: Configuring Multitopology Routing Based on Applications • Example: Configuring Multitopology Routing Based on a Multicast Source

prefix-limit

Syntax	<pre>prefix-limit { maximum <i>number</i>; teardown <<i>percentage</i>> <idle-timeout (forever <i>minutes</i>)>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</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> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit protocols bgp group <i>group-name</i> family (inet inet6) (any labeled-unicast multicast unicast)],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6) (any flow labeled-unicast multicast unicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Limit the number of prefixes received on a BGP peer session and a rate-limit logging when injected prefixes exceed a set limit.
Options	<p>maximum <i>number</i>—When you set the maximum number of prefixes, a message is logged when that number is exceeded.</p> <p>Range: 1 through 4,294,967,295 ($2^{32} - 1$)</p> <p>teardown <<i>percentage</i>>—If you include the teardown statement, the session is torn down when the maximum number of prefixes is reached. If you specify a percentage, messages are logged when the number of prefixes exceeds that percentage. After the session is torn down, it is reestablished in a short time unless you include the idle-timeout statement. Then the session can be kept down for a specified amount of time, or forever. If you specify forever, the session is reestablished only after you issue a clear bgp neighbor command.</p> <p>Range: 1 through 100</p>

idle-timeout (**forever** | *timeout-in-minutes*)—(Optional) If you include the **idle-timeout** statement, the session is torn down for a specified amount of time, or forever. If you specify a period of time, the session is allowed to reestablish after this timeout period. If you specify **forever**, the session is reestablished only after you intervene with a **clear bgp neighbor** command.

Range: 1 through 2400

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

Related Documentation	<ul style="list-style-type: none">• accepted-prefix-limit• Understanding Multiprotocol BGP
------------------------------	---

priority (Protocols IS-IS)

Syntax	<code>priority <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure the interface's priority for becoming the designated router. The interface with the highest priority value becomes that level's designated router.</p> <p>The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.</p> <p>A routing device advertises its priority to become a designated router in its hello packets. On all multiaccess networks, IS-IS uses the advertised priorities to elect a designated router for the network. This routing device is responsible for sending network link-state advertisements, which describe all the routing devices attached to the network. These advertisements are flooded throughout a single area.</p> <p>A routing device's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127. Routing devices with a higher value are more likely to become the designated router.</p>
Options	<i>number</i> —Priority value. Range: 0 through 127 Default: 64
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring IS-IS Designated Routers

priority (Protocols OSPF)

Syntax	<code>priority <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</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)) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Specify the routing device's priority for becoming the designated routing device. The routing device that has the highest priority value on the logical IP network or subnet becomes the network's designated router. You must configure at least one routing device on each logical IP network or subnet to be the designated router. You also should specify a routing device's priority for becoming the designated router on point-to-point interfaces.
Options	<p>number—Routing device's priority for becoming the designated router. A priority value of 0 means that the routing device never becomes the designated router. A value of 1 means that the routing device has the least chance of becoming a designated router.</p> <p>Range: 0 through 255</p> <p>Default: 128</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Designated Router Overview • Example: Controlling OSPF Designated Router Election

qualified-next-hop (Static Routes)

Syntax `qualified-next-hop (address | interface-name) {
 bfd-liveness-detection {
 authentication {
 algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 | meticulous-keyed-sha-1 |
 simple-password);
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 holddown-interval milliseconds;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (1 | automatic);
 }
 interface interface-name;
 metric metric;
 preference preference;
}`

Hierarchy Level `[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options static route destination-prefix],
 [edit logical-systems logical-system-name routing-options rib inet6.0 static route
 destination-prefix],
 [edit logical-systems logical-system-name routing-options static route destination-prefix],
 [edit routing-instances routing-instance-name routing-options static route destination-prefix],
 [edit routing-options rib inet6.0 static route destination-prefix],
 [edit routing-options static route destination-prefix]`

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.
 Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure a static route with multiple possible next hops, each of which can have its own preference value, IGP metric that is used when the route is exported into an IGP, and Bidirectional Forwarding Detection (BFD) settings. If multiple links are operational, the one with the most preferred next hop is used. The most preferred next hop is the one with the lowest preference value.

Options *address*—IPv4, IPv6, or ISO network address of the next hop.

interface-name—Name of the interface on which to configure an independent metric or preference for a static route. To configure an unnumbered interface as the next-hop

interface for a static route, specify **qualified-next-hop *interface-name***, where *interface-name* is the name of the IPv4 or IPv6 unnumbered interface.



NOTE: For an Ethernet interface to be configured as the qualified next hop for a static route, it must be an unnumbered interface.

To configure an Ethernet interface as an unnumbered interface, configure the **unnumbered-address <interface-name>** statement at the [edit interfaces <interface-name> unit <logical-unit-number> family <family-name>] hierarchy level as described in Configuring an Unnumbered Interface.

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Static Route Preferences and Qualified Next Hops • Example: Enabling BFD on Qualified Next Hops in Static Routes

readvertise

Syntax	(readvertise no-readvertise);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-options static (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure whether static routes are eligible to be readvertised by routing protocols:
Default	Static routes are eligible to be readvertised (that is, exported from the routing table into dynamic routing protocols) if a policy to do so is configured. To mark an IPv4 static route as being ineligible for readvertisement, include the no-readvertise statement.
Options	readvertise —Readvertise static routes. Include the readvertise statement when configuring an individual route in the route portion of the static statement to override a no-readvertise option specified in the defaults portion of the statement. no-readvertise —Mark a static route as being ineligible for readvertisement. Include the no-readvertise option when configuring the route.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Controlling Static Routes in Routing and Forwarding Tables• static on page 376

realm

Syntax	<pre> realm (ipv4-unicast ipv4-multicast ipv6-unicast) { area <i>area-id</i> { interface <i>interface-name</i>; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf3],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3],</p> <p>[edit protocols ospf3],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Configure OSPFv3 to advertise address families other than unicast IPv6. Junos OS maps each address family you configure to a separate realm with its own set of neighbors and link-state database.</p>
Options	<p>ipv4-unicast—Configure a realm for IPv4 unicast routes.</p> <p>ipv4-multicast—Configure a realm for IPv4 multicast routes.</p> <p>ipv6-multicast—Configure a realm for IPv6 multicast routes.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Multiple Address Families for OSPFv3

receive (Protocols RIP)

Syntax	<code>receive receive-options;</code>
Hierarchy Level	<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>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure RIP receive options.
Options	<i>receive-options</i> —One of the following: <ul style="list-style-type: none">• both—Accept both RIP version 1 and version 2 packets.• none—Do not receive RIP packets.• version-1—Accept only RIP version 1 packets.• version-2—Accept only RIP version 2 packets. Default: both
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets• send on page 365

receive (Protocols RIPng)

Syntax	receive <none>;
Hierarchy Level	<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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>
Description	Enable or disable receiving of update messages.
Options	<p>none—(Optional) Disable receiving update messages.</p> <p>Default: Enabled</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • send on page 366 • Example: Configuring RIPng


redundant-sources

Syntax	<code>redundant-sources [<i>addresses</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-options multicast flow-map <i>flow-map-name</i>]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure a list of redundant sources for multicast flows defined by a flow map.
Options	<i>addresses</i> —List of IPv4 or IPv6 addresses for use as redundant (backup) sources for multicast flows defined by a flow map.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring a Multicast Flow Map

reference-bandwidth (Protocols IS-IS)

Syntax	<code>reference-bandwidth <i>reference-bandwidth</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],</p> <p>[edit protocols <i>isis</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Optimize routing based on bandwidth by setting the reference bandwidth used in calculating the default interface cost.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. When there are several equal-cost routes to a destination, traffic is distributed equally among them.</p> <p>The cost of a route is described by a single dimensionless metric that is determined using the following formula:</p> $\text{cost} = \text{reference-bandwidth} / \text{bandwidth}$ <p>For example, if you set the reference bandwidth to 1 Gbps (that is, <i>reference-bandwidth</i> is set to 1,000,000,000), a 100-Mbps interface has a routing metric of 10.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics.</p>
Options	<p><i>reference-bandwidth</i>—Reference bandwidth value in bits per second.</p> <p>Range: 9600 through 1,000,000,000,000 bps</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring IS-IS http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

reference-bandwidth (Protocols OSPF)

Syntax	<code>reference-bandwidth <i>reference-bandwidth</i>;</code>
Hierarchy Level	<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> 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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm 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>
Description	<p>Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula:</p> $\text{cost} = \text{ref-bandwidth} / \text{bandwidth}$
Options	<p><i>reference-bandwidth</i>—Reference bandwidth, in bits per second.</p> <p>Range: 9600 through 1,000,000,000,000 bits</p> <p>Default: 100 Mbps (100,000,000 bits)</p>
<div>  <p>NOTE: The default behavior is to use the reference-bandwidth value to calculate the cost of OSPF interfaces. You can override this behavior for any OSPF interface by configuring a specific cost with the metric statement.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Controlling the Cost of Individual OSPF Network Segments metric on page 249

register-probe-time

Syntax	<code>register-probe-time <i>register-probe-time</i>;</code>
Hierarchy Level	[edit protocols pim rp]
Release Information	Statement introduced in Junos OS Release 12.2 for EX Series switches.
Description	Specify the amount of time before the register suppression time (RST) expires when a designated switch can send a NULL-Register to the rendezvous point (RP).
Options	<i>register-probe-time</i> —Amount of time before the RST expires. Default: 5 seconds Range: 5 to 60 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• PIM Overview• Understanding PIM Sparse Mode

remove-private

Syntax `remove-private all replace nearest;`

Hierarchy Level `[edit logical-systems logical-system-name protocols bgp],`
 `[edit logical-systems logical-system-name protocols bgp group group-name],`
 `[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],`
 `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`
 `bgp],`
 `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`
 `bgp group group-name],`
 `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`
 `bgp group group-name neighbor address],`
 `[edit protocols bgp],`
 `[edit protocols bgp group group-name],`
 `[edit protocols bgp group group-name neighbor address],`
 `[edit routing-instances routing-instance-name protocols bgp],`
 `[edit routing-instances routing-instance-name protocols bgp group group-name],`
 `[edit routing-instances routing-instance-name protocols bgp group group-name neighbor`
 `address]`

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description When advertising AS paths to remote systems, have the local system strip private AS numbers from the AS path. The numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. If the AS path contains the AS number of the external BGP (EBGP) neighbor, BGP does not remove the private AS number.



NOTE: As of Junos OS 10.0R2 and higher, if there is a need to send prefixes to an EBGP peer that has an AS number that matches an AS number in the AS path, consider using the `as-override` statement instead of the `remove-private` statement.

The operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.

The Junos OS recognizes the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document.

The set of reserved AS numbers is in the range from 64,512 through 65,535.

Options **all**—Remove all private AS numbers from the original path. Do not stop the process of removing private AS numbers, even if a public AS number is encountered.


nearest—When you use the **all** and **replace** options, choose the last (right-most) public AS number encountered in the original AS path for the replacement value, as the AS path is processed from left to right. If no public AS number is encountered, the default replacement value is used. (See the **replace** option for information about the default replacement value.)

replace—When you use the **all** option, instead of removing private AS numbers, perform a replace operation. The default replacement value for the private AS number is the local AS number at the BGP group level for the BGP peer. If you are unsure about the replacement value, check the local AS value displayed in the output of the **show bgp group group-name** command.

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

Related Documentation	<ul style="list-style-type: none">• Example: Removing Private AS Numbers from AS Paths
------------------------------	--

resolution

Syntax	<pre>resolution { rib routing-table-name { import [policy-names]; resolution-ribs [routing-table-names]; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Configure the router to perform custom route resolution on protocol next hops of routes in a certain routing table. The protocol next hop is used to determine the forwarding next-hop.</p> <p>For example, you might want to direct inet.2 route resolution to use topology routing tables :red.inet.0 and :blue.inet.0 for protocol next-hop IP address lookups. Or you might want to direct bgp.l3vpn.0 to use the information in inet.0 to resolve routes, thus overriding the default behavior, which is to use inet.3.</p> <p>You can specify up to two routing tables in the resolution-ribs statement. The route resolution scheme first checks the first-listed routing table for the protocol next-hop address. If the address is found, it uses this entry. If it is not found, the resolution scheme checks second-listed routing table. Hence, only one routing table is used for each protocol nexthop address. For example, if you configure resolution rib bgp.l3vpn.0 resolution-ribs [inet.0 inet.3], inet.0 is checked first and then inet.3 is checked.</p> <div><p>NOTE: Customizing route resolution might cause the routing protocol process (rpd) to consume more memory resources than it ordinarily would. When you customize route resolution, we recommend that you check the memory resources by running the show system processes and the show task memory commands. For more information, see Routing Protocol Process Memory FAQs.</p></div>
The remaining statements are explained separately.	
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

- Related Documentation**
- Example: Configuring Route Resolution on PE Routers
 - Example: Configuring Route Resolution on Route Reflectors
 - Example: Configuring Multitopology Routing Based on a Multicast Source

resolution-ribs

Syntax	<code>resolution-ribs [<i>routing-table-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution rib],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options resolution rib],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options resolution rib],</p> <p>[edit routing-options resolution rib]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify one or more routing tables to use for route resolution.</p> <p>This statement enables you to override the default routing tables that Junos OS uses for route resolution. For example, suppose that the resolution routing table is inet.3, but you want to allow fallback resolution through inet.0. One example use case is overriding the bgp.rtarget.0 (family route-target) routing table resolution from using only inet.3 to using both inet.3 and inet.0.</p>
Options	<i>routing-table-names</i> —Name of one or more routing tables.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Route Resolution on PE Routers • Example: Configuring Multitopology Routing Based on a Multicast Source

resolve

Syntax	resolve;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-options static (defaults route)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Statically configure routes to be resolved to a next hop that is not directly connected. The route is resolved through the inet.0 and inet.3 routing tables.
Default	Static routes can point only to a directly connected next hop.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• static on page 376

restart-duration (Routing Options)

Syntax	<code>restart-duration <i>seconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart],</p> <p>[edit routing-options graceful-restart]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Configure the restart timer for graceful restart.
Options	<p>seconds—Configure the time period for the restart to last.</p> <p>Range: 120 through 900 seconds</p> <p>Default: 300 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Junos OS High Availability Configuration Guide

retain

Syntax	(no-retain retain);
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)], [edit routing-options static (defaults route)]</pre>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure statically configured routes to be deleted from or retained in the forwarding table when the routing protocol process shuts down normally:
Default	Statically configured routes are deleted from the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.
Options	<p>no-retain—Delete statically configured routes from the forwarding table when the routing protocol process shuts down normally. To explicitly specify that routes be deleted from the forwarding table, include the no-retain statement. Include this statement when configuring an individual route in the route portion of the static statement to override a retain option specified in the defaults portion of the statement.</p> <p>retain—Have a static route remain in the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Examples: Configuring Static Routes• static on page 376

retransmit-interval (OSPF)

Syntax	<code>retransmit-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</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) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Specify how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements (LSAs) to an interface's neighbors.
Options	<p>seconds—Interval to wait.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 5 seconds</p>



NOTE: You must configure LSA retransmit intervals to be equal to or greater than 3 seconds to avoid triggering a retransmit trap, because Junos OS delays LSA acknowledgments by up to 2 seconds.

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

Related Documentation

- Example: Configuring OSPF Timers
- Configuring RSVP and OSPF for LMP Peer Interfaces

reverse-oif-mapping

Syntax `reverse-oif-mapping {
 no-qos-adjust;
}`

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast [interface interface-name](#)],
[edit logical-systems *logical-system-name* routing-options multicast [interface interface-name](#)],
[edit routing-instances *routing-instance-name* routing-options multicast [interface interface-name](#)],
[edit routing-options multicast [interface interface-name](#)]

Release Information Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 9.2 for EX Series switches.
The **no-qos-adjust** statement added in Junos OS Release 9.5.
The **no-qos-adjust** statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Enable the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD request it receives over the multicast VLAN.

The remaining statement is explained separately.

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

Related Documentation

- Example: Configuring Multicast with Subscriber VLANs

rib (General)

```
Syntax  rib routing-table-name {
        aggregate {
            defaults {
                ... aggregate-options ...
            }
            route destination-prefix {
                policy policy-name;
                ... aggregate-options ...
            }
        }
        generate {
            defaults {
                generate-options;
            }
            route destination-prefix {
                policy policy-name;
                generate-options;
            }
        }
        martians {
            destination-prefix match-type <allow>;
        }
    }
    static {
        defaults {
            static-options;
        }
        rib-group group-name;
        route destination-prefix {
            next-hop;
            static-options;
        }
    }
}
```

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Create a routing table.

Explicitly creating a routing table with ***routing-table-name*** is optional if you are not adding any static, martian, aggregate, or generated routes to the routing table and if you also are creating a routing table group.



NOTE: The IPv4 multicast routing table (inet.1) and the IPv6 multicast routing table (inet6.1) are not supported for this statement.

Default If you do not specify a routing table name with the *routing-table-name* option, the software uses the default routing tables, which are **inet.0** for unicast routes and **inet.1** for the multicast cache.

Options *routing-table-name*—Name of the routing table, in the following format:
protocol [.identifier].

In a routing instance, the routing table name must include the routing instance name.

For example, if the routing instance name is **link0**, the routing table name might be **link0.inet6.0**.

- *protocol* is the protocol family. It can be **inet6** for the IPv6 family, **inet** for the IPv4 family, **iso** for the ISO protocol family, or *instance-name.iso.0* for an ISO routing instance.
- *identifier* is a positive integer that specifies the instance of the routing table.

Default: **inet.0**

The remaining statements are explained separately.

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

Related Documentation

- Example: Creating Routing Tables
- passive

rib (Route Resolution)

Syntax	<pre> rib <i>routing-table-name</i> { import [<i>policy-names</i>]; resolution-ribs [<i>routing-table-names</i>]; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options resolution],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options resolution],</p> <p>[edit routing-options resolution]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify a routing table name for route resolution.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Route Resolution on PE Routers

rib-group (Protocols BGP)

Syntax	<code>rib-group group-name;</code>
Hierarchy Level	<code>[edit logical-systems logical-system-name protocols bgp family inet (labeled-unicast unicast multicast)],</code> <code>[edit logical-systems logical-system-name protocols bgp group group-name family inet (labeled-unicast unicast multicast)],</code> <code>[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family inet (labeled-unicast unicast multicast)],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp family inet (labeled-unicast unicast multicast)],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name family inet (labeled-unicast unicast multicast)],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name neighbor address family inet (labeled-unicast unicast multicast)],</code> <code>[edit protocols bgp family inet (labeled-unicast unicast multicast)],</code> <code>[edit protocols bgp group group-name family inet (labeled-unicast unicast multicast)],</code> <code>[edit protocols bgp group group-name neighbor address family inet (labeled-unicast unicast multicast)],</code> <code>[edit routing-instances routing-instance-name protocols bgp family inet (labeled-unicast unicast multicast)],</code> <code>[edit routing-instances routing-instance-name protocols bgp group group-name family inet (labeled-unicast unicast multicast)],</code> <code>[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address family inet (labeled-unicast unicast multicast)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Add unicast prefixes to unicast and multicast tables.
Options	group-name —Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. You generally specify only one routing table group.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Exporting Specific Routes from One Routing Table Into Another Routing Table• Example: Importing Direct and Static Routes Into a Routing Instance• Understanding Multiprotocol BGP

rib-group (Protocols IS-IS)

Syntax	<pre> rib-group { inet <i>group-name</i>; inet6 <i>group-name</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Install routes learned from IS-IS routing instances into routing tables in the IS-IS routing table group. You can install IPv4 routes or IPv6 routes.</p> <p>Support for IPv6 routing table groups in IS-IS enables IPv6 routes that are learned from IS-IS routing instances to be installed into other routing tables defined in an IS-IS routing table group.</p>
Options	<p><i>group-name</i>—Name of the routing table group.</p> <p>inet—Install IPv4 IS-IS routes.</p> <p>inet6—Install IPv6 IS-IS routes.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Exporting Specific Routes from One Routing Table Into Another Routing Table • Example: Importing Direct and Static Routes Into a Routing Instance • Understanding Multiprotocol BGP

rib-group (Protocols OSPF)

Syntax	<code>rib-group <i>group-name</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> (ospf ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit protocols (ospf ospf3)],</code> <code>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Install routes learned from OSPF routing instances into routing tables in the OSPF routing table group.
Options	<i>group-name</i> —Name of the routing table group.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Exporting Specific Routes from One Routing Table Into Another Routing Table• Example: Importing Direct and Static Routes Into a Routing Instance• Understanding Multiprotocol BGP• interface-routes on page 212• rib-group on page 352

rib-group (Protocols RIP)

Syntax	<code>rib-group <i>group-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rip], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip], [edit protocols rip], [edit routing-instances <i>routing-instance-name</i> protocols rip]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Install RIP routes into multiple routing tables by configuring a routing table group.
Options	<i>group-name</i> —Name of the routing table group.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Redistributing Routes Between Two RIP Instances

rib-group (Routing Options)

Syntax	<code>rib-group <i>group-name</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options interface-routes],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options interface-routes],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options static],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options interface-routes],</code> <code>[edit routing-options interface-routes],</code> <code>[edit routing-options rib <i>routing-table-name</i> static],</code> <code>[edit routing-options static]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure which routing table groups interface routes are imported into.
Options	<i>group-name</i> —Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. It generally does not make sense to specify more than a single routing table group.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Importing Direct and Static Routes Into a Routing Instance• Example: Exporting Specific Routes from One Routing Table Into Another Routing Table• interface-routes on page 212• rib-groups on page 353

rib-groups

Syntax	<pre> rib-groups { group-name { export-rib group-name; import-policy [policy-names]; import-rib [group-names]; } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Group one or more routing tables to form a routing table group. A routing protocol can import routes into all the routing tables in the group and can export routes from a single routing table.</p> <p>Each routing table group must contain one or more routing tables that Junos OS uses when importing routes (specified in the import-rib statement) and optionally can contain one routing table group that Junos OS uses when exporting routes to the routing protocols (specified in the export-rib statement).</p> <p>The first routing table you specify is the <i>primary routing table</i>, and any additional routing tables are the <i>secondary routing tables</i>.</p> <p>The primary routing table determines the address family of the routing table group. To configure an IP version 4 (IPv4) routing table group, specify inet.0 as the primary routing table. To configure an IP version 6 (IPv6) routing table group, specify inet6.0 as the primary routing table. If you configure an IPv6 routing table group, the primary and all secondary routing tables must be IPv6 routing tables (inet6.x).</p> <p>In Junos OS Release 9.5 and later, you can include both IPv4 and IPv6 routing tables in an IPv4 import routing table group using the import-rib statement. In releases prior to Junos OS Release 9.5, you can only include either IPv4 or IPv6 routing tables in the same import-rib statement. The ability to configure an import routing table group with both IPv4 and IPv6 routing tables enables you, for example, to populate the inet6.3 routing table with IPv6 addresses that are compatible with IPv4. Specify inet.0 as the primary routing table, and specify inet6.3 as a secondary routing table.</p>



NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.



NOTE: If you configure an import routing table group that includes both IPv4 and IPv6 routing tables, any corresponding export routing table group must include only IPv4 routing tables.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group. For more information, see Example: Configuring Multiple Routing Instances of OSPF.

After specifying the routing table from which to import routes, you can apply one or more policies to control which routes are installed in the routing table group. To apply a policy to routes being imported into the routing table group, include the **import-policy** statement.

Options *group-name*—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens.

The remaining statements are explained separately.

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

Related Documentation

- Example: Exporting Specific Routes from One Routing Table Into Another Routing Table
- [rib-group on page 352](#)


rip

Syntax	rip {...}
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable RIP routing on the routing device.
Default	RIP is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Configuring RIP

ripng

Syntax	ripng {...}
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for routing instances introduced in Junos OS Release 9.0.
Description	Enable RIPng routing on the routing device.
Default	RIPng is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Configuring RIPng

route-distinguisher-id

Syntax	<code>route-distinguisher-id <i>ip-address</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Automatically assign a route distinguisher to the routing instance.</p> <p>If you configure the route-distinguisher statement in addition to the route-distinguisher-id statement, the value configured for route-distinguisher supersedes the value generated from route-distinguisher-id.</p> <div><p>NOTE: To avoid a conflict in the two route distinguisher values, it is recommended to ensure that the first half of the route distinguisher obtained by configuring the route-distinguisher statement is different from the first half of the route distinguisher obtained by configuring the route-distinguisher-id statement.</p></div>
Options	<i>ip-address</i> —Address for routing instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BGP Route Target Filtering for VPNs• Configuring Routing Instances on PE Routers in VPNs

route-record

Syntax	route-record;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Export the AS path and routing information to the traffic sampling process. Before you can perform flow aggregation, the routing protocol process must export the AS path and routing information to the sampling process.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling Flow Aggregation• Junos Services Interfaces Configuration Release 11.2

route-timeout (Protocols RIP)

Syntax	<code>route-timeout seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols rip],</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> protocols rip group <i>group-name</i>],</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 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> 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>],</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>],</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>
Release Information	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the route timeout interval for RIP. If a route is not refreshed after being installed in the routing table by the specified timeout interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.
Options	seconds —Estimated time to wait before making updates to the routing table. Range: 30 through 360 seconds Default: 180 seconds
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring RIP Timers• RIP Demand Circuits Overview


route-timeout (Protocols RIPng)

Syntax	<code>route-timeout <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ripng], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng], [edit protocols ripng], [edit routing-instances <i>routing-instance-name</i> protocols ripng]
Release Information	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for routing instances introduced in Junos OS Release 9.0.
Description	Configure the route timeout interval for RIPng.
Options	<i>seconds</i> —Estimated time to wait before making updates to the routing table. Range: 30 through 360 seconds Default: 180 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring RIPng Timers

route-type-community

Syntax	route-type-community (iana vendor);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specify an extended community value to encode the OSPF route type. Each extended community is coded as an eight-octet value. This statement sets the most significant bit to either an IANA or vendor-specific route type.
Options	iana —Encode a route type with the value 0x0306 . This is the default value. vendor —Encode the route type with the value 0x8000 .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring Routing Between PE and CE Routers in Layer 3 VPNs

router-id

Syntax	<code>router-id address;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Specify the routing device's IP address.</p> <p>The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier usually is the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;">  <p>NOTE: We strongly recommend that you configure the router identifier under the [edit routing-options] hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.</p> </div> <p>For more information about the router identifier in OSPF, see Example: Configuring an OSPF Router Identifier.</p>
Options	<p>address—IP address of the routing device.</p> <p>Default: Address of the first interface encountered by Junos OS</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Examples: Configuring External BGP Peering Examples: Configuring Internal BGP Peering

routing-options

Syntax	routing-options { ... }
Hierarchy Level	[edit], [edit logical-systems <i>logical-system-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure protocol-independent routing properties.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Protocol-Independent Routing Properties Configuration Guide

rpf-check-policy (Routing Options RPF)

Syntax	rpf-check-policy [<i>policy-names</i>];
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 8.1. 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.
Description	Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.
Options	<i>policy-names</i> —Name of one or more multicast RPF check policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring RPF Policies

scope

Syntax	<pre>scope scope-name { interface [interface-names]; prefix destination-prefix; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Configure multicast scoping.
Options	<p><i>scope-name</i>—Name of the multicast scope.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Creating a Named Scope for Multicast Scoping

scope-policy

Syntax `scope-policy [policy-names];`

Hierarchy Level `[edit logical-systems logical-system-name routing-options multicast],`
`[edit routing-options multicast]`



NOTE: You can configure a scope policy at these two hierarchy levels only. You cannot apply a scope policy to a specific routing instance, because all scoping policies are applied to all routing instances. However, you can apply the `scope` statement to a specific routing instance at the `[edit routing-instances routing-instance-name routing-options multicast]` or `[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options multicast]` hierarchy level.

Release Information Statement introduced before Junos 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.

Description Apply policies for scoping. The policy must be correctly configured at the **edit policy-options policy-statement** hierarchy level.

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

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

Related Documentation

- [scope on page 363](#)
- Example: Using a Scope Policy for Multicast Scoping

send (Protocols RIP)

Syntax	<code>send <i>send-options</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rip],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip 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 rip],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols rip],</p> <p>[edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure RIP send options.
Options	<p><i>send-options</i>—One of the following:</p> <ul style="list-style-type: none"> • broadcast—Broadcast RIP version 2 packets (RIP version 1 compatible). • multicast—Multicast RIP version 2 packets. This is the default. • none—Do not send RIP updates. • version-1—Broadcast RIP version 1 packets. <p>Default: multicast</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets • receive on page 330

send (Protocols RIPng)

Syntax	send <none>;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ripng], [edit logical-systems <i>logical-system-name</i> protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances-name</i> protocols ripng], [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>], [edit protocols ripng], [edit protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ripng], [edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i>]
Release Information	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.
Description	Enable or disable sending of update messages.
Options	none —(Optional) Disable sending of update messages. Default: Enabled
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• receive on page 331


shortcuts (Protocols OSPF)

Syntax	shortcuts { lsp-metric-into-summary; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) traffic-engineering], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) traffic-engineering], [edit protocols (ospf ospf3) traffic-engineering], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) traffic-engineering]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4. Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4 for EX Series switches.
Description	Configure OSPF to use MPLS label-switched paths (LSPs) as shortcut next hops. By default, shortcut routes calculated through OSPFv2 are installed in the inet.3 routing table, and shortcut routes calculated through OSPFv3 are installed in the inet6.3 routing table.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Enabling OSPF Traffic Engineering Support

source (Source-Specific Multicast)

Syntax	source [<i>addresses</i>];
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast ssm-map <i>ssm-map-name</i>], [edit routing-options multicast ssm-map <i>ssm-map-name</i>]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specify IPv4 or IPv6 source addresses for an SSM map.
Options	addresses —IPv4 or IPv6 source addresses.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To view this statement in the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring SSM Mapping

source-routing

Syntax	source-routing { (ip ipv6) }
Hierarchy Level	[edit routing-options]
Release Information	Statement for IPv6 introduced in Junos OS Release 8.2. Statement for IPv4 introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Enable source routing.</p> <p>Source routing allows a sender of a packet to partially or completely specify the route the packet takes through the network. In contrast, in non-source routing protocols, routers in the network determine the path based on the packet's destination.</p> <div>  <p>NOTE: We recommend that you not use source routing. Instead, we recommend that you use policy-based routing or filter-based forwarding to route packets based on source addresses.</p> </div>
Default	Disabled
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Filter-Based Forwarding on the Source Address

spf-options (Protocols IS-IS)

Syntax	<pre>spf-options { delay <i>milliseconds</i>; holddown <i>milliseconds</i>; rapid-runs <i>number</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>], [edit protocols <i>isis</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure options for running the shortest-path-first (SPF) algorithm.</p> <p>Running the SPF algorithm is usually the beginning of a series of larger system-wide events. For example, the SPF algorithm can lead to interior gateway protocol (IGP) prefix changes, which then lead to BGP nexthop resolution changes. Consider what happens if there are rapid link changes in the network. The local routing device can become overwhelmed. This is why it sometimes makes sense to throttle the scheduling of the SPF algorithm.</p> <p>You can configure the following SPF options:</p> <ul style="list-style-type: none">• The delay in the time between the detection of a topology change and when the SPF algorithm actually runs.• The maximum number of times that the SPF algorithm can run in succession before the hold-down timer begins.• The time to hold down, or wait, before running another SPF calculation after the SPF algorithm has run in succession the configured maximum number of times. <p>If the network stabilizes during the hold-down period and the SPF algorithm does not need to run again, the system reverts to the configured values for the delay and rapid-runs statements.</p>
Options	<p>delay <i>milliseconds</i>—Time interval between the detection of a topology change and when the SPF algorithm runs.</p> <p>Range: 50 through 1000 milliseconds</p> <p>Default: 200 milliseconds</p> <p>holddown <i>milliseconds</i>—Time interval to hold down, or wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession.</p> <p>Range: 2000 through 10,000 milliseconds</p> <p>Default: 5000 milliseconds</p>

rapid-runs *number*—Maximum number of times the SPF algorithm can run in succession.
After the maximum is reached, the holddown interval begins.

Range: 1 through 5

Default: 3

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

Related Documentation	<ul style="list-style-type: none">• Example: Configuring Link and Node Protection for IS-IS Routes
------------------------------	--

spf-options (Protocols OSPF)

Syntax	<pre>spf-options { delay <i>milliseconds</i>; holddown <i>milliseconds</i>; rapid-runs <i>number</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (ospf ospf3)], [edit protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm 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>
Description	<p>Configure options for running the shortest-path-first (SPF) algorithm. You can configure the following:</p> <ul style="list-style-type: none">• A delay for when to run the SPF algorithm after a network topology change is detected.• The maximum number of times the SPF algorithm can run in succession.• A hold-down interval after the SPF algorithm runs the maximum number of times. <p>Running the SPF algorithm is usually the beginning of a series of larger system-wide events. For example, the SPF algorithm can lead to interior gateway protocol (IGP) prefix changes, which then lead to BGP nexthop resolution changes. Consider what happens if there are rapid link changes in the network. The local routing device can become overwhelmed. This is why it sometimes makes sense to throttle the scheduling of the SPF algorithm.</p>

Options	delay <i>milliseconds</i> —Time interval between the detection of a topology change and when the SPF algorithm runs. Range: 50 through 8000 milliseconds Default: 200 milliseconds
	holddown <i>milliseconds</i> —Time interval to hold down, or to wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession. Range: 2000 through 20,000 milliseconds Default: 5000 milliseconds
	rapid-runs <i>number</i> —Maximum number of times the SPF algorithm can run in succession. After the maximum is reached, the hold down interval begins. Range: 1 through 10 Default: 3
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring SPF Algorithm Options for OSPF• Example: Configuring Multitopology Routing Based on Applications• Example: Configuring Multitopology Routing Based on a Multicast Source

ssm-groups

Syntax	<code>ssm-groups [<i>ip-addresses</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</code> <code>[edit routing-options multicast]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Configure source-specific multicast (SSM) groups.</p> <p>By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the ssm-groups statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the ssm-groups statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.</p> <p>IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.</p>
Options	<i>ip-addresses</i> —List of one or more additional SSM group addresses separated by a space.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring Source-Specific Multicast Groups with Any-Source Override

ssm-map (Routing Options Multicast)

Syntax	<pre>ssm-map <i>ssm-map-name</i> { policy [<i>policy-names</i>]; source [<i>addresses</i>]; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Configure SSM mapping.
Options	<p><i>ssm-map-name</i>—Name of the SSM map.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring SSM Mapping

static (Routing Options)

```
Syntax static {
    defaults {
        static-options;
    }
    rib-group group-name;
    route destination-prefix {
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            local-address ip-address;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-receive-ttl number;
            multiplier number;
            neighbor address;
            no-adaptation;
            transmit-interval {
                threshold milliseconds;
                minimum-interval milliseconds;
            }
            version (1 | automatic);
        }
        next-hop address;
        next-hop options;
        qualified-next-hop address {
            bfd-liveness-detection {
                authentication {
                    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                        meticulous-keyed-sha-1 | simple-password);
                    key-chain key-chain-name;
                    loose-check;
                }
                detection-time {
                    threshold milliseconds;
                }
                holddown-interval milliseconds;
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                version (1 | automatic);
            }
        }
    }
}
```

```

        metric metric;
        preference preference;
    }
    static-options;
}

```

Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options],</p> <p>[edit routing-options rib <i>routing-table-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos 9.6.</p> <p>Support for BFD authentication introduced in Junos 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Configure static routes to be installed in the routing table. You can specify any number of routes within a single static statement, and you can specify any number of static options in the configuration.</p>

Options defaults—(Optional) Specify global static route options. These options only set default attributes inherited by all newly created static routes. These are treated as global defaults and apply to all the static routes you configure in the **static** statement.



NOTE: Specifying the global static route options does not create default routes. These options only set default attributes inherited by all newly created static routes.

route—Configure individual static routes. In this part of the **static** statement, you optionally can configure static route options. These options apply to the individual destination only and override any options you configured in the **defaults** part of the **static** statement.

- **destination-prefix/prefix-length—destination-prefix** is the network portion of the IP address, and **prefix-length** is the destination prefix length.

When you configure an individual static route in the **route** part of the **static** statement, specify the destination of the route (in **route destination-prefix**) in one of the following ways:

- **network/mask-length**, where **network** is the network portion of the IP address and **mask-length** is the destination prefix length.
- **default** if this is the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.



NOTE: IPv4 packets with a destination of 0.0.0.0 (the obsoleted limited broadcast address) and IPv6 packets with a destination of 0::0 are discarded by default. To forward traffic destined to these addresses, you can add a static route to 0.0.0.0/32 for IPv4 or 0::0/128 for IPv6.

- **nsap-prefix—nsap-prefix** is the network service access point (NSAP) address for ISO.
- **next-hop address**—Reach the next-hop routing device by specifying an IP address, an interface name, or an ISO network entity title (NET).

IPv4 or IPv6 address of the next hop to the destination, specified as:

- IPv4 or IPv6 address of the next hop
- Interface name (for point-to-point interfaces only)
- **address** or **interface-name** to specify an IP address of a multipoint interface or an interface name of a point-to-point interface.



NOTE: If an interface becomes unavailable, all configured static routes on that interface are withdrawn from the routing table.



NOTE: 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.

next-hop options—Additional information for how to manage forwarding of packets to the next hop.

- **discard**—Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP (or ICMPv6) unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.
- **iso-net**—Reach the next-hop routing device by specifying an ISO NSAP.

- **next-table *routing-table-name***—Name of the next routing table to the destination.

If you use the **next-table** action, the configuration must include a term qualifier that specifies a different table than the one specified in the **next-table** action. In other words, the term qualifier in the **from** statement must exclude the table in the **next-table** action. In the following example, the first term contains **rib vrf-customer2.inet.0** as a matching condition. The action specifies a next-hop in a different routing table, **vrf-customer1.inet.0**. The second term does the opposite by using **rib vrf-customer1.inet.0** in the match condition and **vrf-customer2.inet.0** in the **next-table** action.

```
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;  
  }  
}
```



NOTE: Within a routing instance, you cannot configure a static route with the **next-table inet.0** statement if any static route in the main routing instance is already configured with the **next-table** statement to point to the **inet.0** routing table of the routing instance. For example, if you configure on the main routing instance a static route 192.168.88.88/32 with the **next-table test.inet.0** statement and the routing instance **test** is also configured with a static route 192.168.88.88/32 with the **next-table inet.0** statement, the commit operation fails. Instead, you must configure a routing table group both on the main instance and on the routing instance, which enables you to install the static route into both routing tables.

- **receive**—Install a route for this next-hop destination into the routing table.

The **receive** option forces the packet to be sent to the Routing Engine.

The **receive** option can be useful in the following cases:

- For receiving MPLS packets destined to a VRF instance's loopback address
- For receiving packets on a link's subnet address, with zeros in the host portion of the address
- **reject**—Do not forward packets addressed to this destination. Instead, drop the packets, send ICMP (or ICMPv6) unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

static-options—(Optional under **route**) Additional information about static routes, which is included with the route when it is installed in the routing table.

You can specify one or more of the following in **static-options**. Each of the options is explained separately.

- (**active** | **passive**);
- **as-path** <as-path> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator as-number in-address>;
- **community** [*community-ids*];
- (**install** | **no-install**);
- (**metric** | **metric2** | **metric3** | **metric4**) *value* <type type>;
- (**preference** | **preference2** | **color** | **color2**) *preference* <type type>;
- (**readvertise** | **no-readvertise**);
- (**resolve** | **no-resolve**);
- (**retain** | **no-retain**);
- **tag** *string*;

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Examples: Configuring Static Routes

stub

Syntax	stub < default-metric <i>metric</i> > <(no-summaries summaries)>;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i>],</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) area <i>area-id</i>],</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) area <i>area-id</i>],</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) area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Specify that this area not be flooded with AS external link-state advertisements (LSAs). You must include the stub statement when configuring all routing devices that are in the stub area.</p> <p>The backbone cannot be configured as a stub area.</p> <p>You cannot configure an area to be both a stub area and a not-so-stubby area (NSSA).</p>
Options	<p>no-summaries—(Optional) Do not advertise routes into the stub area. If you include the default-metric option, only the default route is advertised.</p> <p>summaries—(Optional) Flood summary LSAs into the stub area.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• OSPF Areas and Router Functionality Overview• Example: Configuring OSPF Stub and Totally Stubby Areas• nssa on page 288

subscriber-leave-timer

Syntax	<code>subscriber-leave-timer seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit routing-options multicast interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>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> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message.
Options	<p>seconds—Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message. Specifying a value of 0 results in an immediate update. This is the same as if the statement were not configured.</p> <p>Range: 0 through 30</p> <p>Default: 0 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Multicast with Subscriber VLANs

summaries

Syntax	(summaries no-summaries);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area area-id nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id nssa</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id nssa],</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) area <i>area-id nssa</i>],</p> <p>[edit protocols (ospf ospf3) area area-id nssa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] area <i>area-id nssa</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id nssa</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Configure whether or not area border routers advertise summary routes into an not-so-stubby area (NSSA):</p> <ul style="list-style-type: none">• summaries—Flood summary link-state advertisements (LSAs) into the NSSA.• no-summaries—Prevent area border routers from advertising summaries into an NSSA. If default-metric is configured for an NSSA, a Type 3 LSA is injected into the area by default.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• OSPF Areas and Router Functionality Overview• Example: Configuring OSPF Not-So-Stubby Areas• nssa on page 288• stub on page 382

tag (Routing Options)

Syntax	<code>tag <i>string</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (aggregate generate static) (defaults route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options aggregate generate static) (defaults route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)],</p> <p>[edit routing-options (aggregate generate static) (defaults route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (aggregate generate static) (defaults route)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Associate an OSPF tag with a static, aggregate, or generated route.
Default	No OSPF tag strings are associated with static routes.
Options	<i>string</i> —OSPF tag string.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Examples: Configuring Static Routes • Example: Summarizing Routes Through Route Aggregation • Example: Conditionally Generating Static Routes • aggregate on page 74 • generate on page 164 • static on page 376

tcp-mss (Protocols BGP)

Syntax	<code>tcp-mss <i>segment-size</i>;</code>
Hierarchy Level	<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</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> 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>neighbor-name</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocol bgp group <i>group-name</i>],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</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> neighbor</code> <code> <i>neighbor-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the maximum segment size (MSS) for the TCP connection for BGP neighbors.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Limiting TCP Segment Size for BGP

threshold (Multicast Forwarding Cache)

Syntax	<pre>threshold { log-warning <i>value</i>; suppress <i>value</i> <reuse <i>value</i>>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache (inet inet6)],</p> <p>[edit routing-options multicast forwarding-cache],</p> <p>[edit routing-options multicast forwarding-cache family (inet inet6)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>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> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Configure the global suppression, reuse, and warning log message thresholds for multicast forwarding cache limits. You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the threshold statement globally for the multicast forwarding cache or including the family statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p> <p>To confirm the configured threshold values, use the show multicast forwarding-cache statistics command.</p>
Options	<p>reuse <i>value</i>—(Optional) Value at which to begin creating new multicast forwarding cache entries. If configured, this number should be less than the suppress value.</p> <p>Range: 1 through 200,000</p> <p>suppress <i>value</i>—Value at which to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number should be greater than the reuse value.</p> <p>Range: 1 through 200,000</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Examples: Configuring the Multicast Forwarding Cache

timeout (Flow Maps)

Syntax	timeout (never non-discard-entry-only <i>minutes</i>);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options multicast flow-map <i>flow-map-name</i>], [edit routing-options multicast flow-map <i>flow-map-name</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure the timeout value for multicast forwarding cache entries associated with the flow map.
Options	minutes —Length of time that the forwarding cache entry remains active. Range: 1 through 720 never non-discard-entry-only —Specify that the forwarding cache entry always remain active. If you omit the non-discard-entry-only option, all multicast forwarding entries, including those in forwarding and pruned states, are kept forever. If you include the non-discard-entry-only option, entries with forwarding states are kept forever, and entries with pruned states time out.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.


timeout (Multicast)

Syntax	<code>timeout <i>minutes</i> <family (inet inet6)>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],</p> <p>[edit routing-options multicast forwarding-cache]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Configure the timeout value for multicast forwarding cache entries.
Options	<p><i>minutes</i>—Length of time that the forwarding cache limit remains active.</p> <p>Range: 1 through 720</p> <p><i>family (inet inet6)</i>—(Optional) Apply the configured timeout to either IPv4 or IPv6 multicast forwarding cache entries. Configuring the timeout statement globally for the multicast forwarding cache or including the family statement to configure the timeout value for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p> <p>Default: By default, the configured timeout applies to both IPv4 and IPv6 multicast forwarding cache entries.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring the Multicast Forwarding Cache

topologies (Protocols IS-IS)

Syntax	<pre>topologies { ipv4-multicast; ipv6-multicast; ipv6-unicast; }</pre>
Hierarchy Level	[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]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure alternate IS-IS topologies. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies• Example: Configuring IS-IS Multicast Topology

traceoptions (Protocols BGP)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [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>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>4byte-as statement introduced in Junos OS Release 9.2.</p> <p>4byte-as statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Configure BGP protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.
<div>  <p>NOTE: The traceoptions statement is not supported on QFabric systems.</p> </div>	
Default	The default BGP protocol-level tracing options are inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level. The default group-level trace options are inherited from the BGP protocol-level traceoptions statement. The default peer-level trace options are inherited from the group-level traceoptions statement.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place BGP tracing output in the file bgp-log.</p>

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

Range: 2 through 1000 files

Default: 10 files

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

BGP Tracing Flags

- **4byte-as**—4-byte AS events.
- **bfd**—BFD protocol events.
- **damping**—Damping operations.
- **graceful-restart**—Graceful restart events.
- **keepalive**—BGP keepalive messages. If you enable the the BGP **update** flag only, received keepalive messages do not generate a trace message.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **open**—Open packets. These packets are sent between peers when they are establishing a connection.
- **packets**—All BGP protocol packets.
- **refresh**—BGP refresh packets.
- **update**—Update packets. These packets provide routing updates to BGP systems. If you enable only this flag, received keepalive messages do not generate a trace message. Use the **keepalive** flag to generate a trace message for keepalive messages.

Global Tracing Flags

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

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

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

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

- **detail**—Provide detailed trace information.
- **filter**—Provide filter trace information. Applies only to **route** and **damping** tracing flags.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-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.

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

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

Default: 128 KB

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

Required Privilege Level	<p>routing and trace—To view this statement in the configuration.</p> <p>routing-control and trace-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • log-updown on page 233 statement • Understanding Trace Operations for BGP Protocol Traffic • Configuring OSPF Refresh and Flooding Reduction in Stable Topologies

traceoptions (Protocols IS-IS)

Syntax	<pre>traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols 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]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure IS-IS protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.



NOTE: The **traceoptions** statement is not supported on QFabric systems.

Default	The default IS-IS protocol-level tracing options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks (" "). All files are placed in the directory /var/log. We recommend that you place IS-IS tracing output in the file isis-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 10 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one flag, include multiple flag statements.</p> <p>IS-IS Protocol-Specific Tracing Flags</p>

- **csn**—Complete sequence number PDU (CSNP) packets
- **error**—Errored IS-IS packets
- **graceful-restart**—Graceful restart operation
- **hello**—Hello packets
- **ldp-synchronization**—Synchronization between IS-IS and LDP
- **lsp**—Link-state PDUs
- **lsp-generation**—Link-state PDU generation packets
- **packets**—All IS-IS protocol packets
- **psn**—Partial sequence number PDU (PSNP) packets
- **spf**—Shortest-path-first calculations

Global Tracing Flags

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

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

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

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

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-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. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

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


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

Default: 128 KB

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

Required Privilege Level	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces• Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces• Example: Enabling Packet Checksums on IS-IS Interfaces

traceoptions (Protocols OSPF)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm 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>
Description	<p>Configure OSPF protocol-level tracing options.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
	<div>  <p>NOTE: The traceoptions statement is not supported on QFabric systems.</p> </div>
Default	The default OSPF protocol-level tracing options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place OSPF tracing output in the file ospf-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p>

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

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

OSPF Tracing Flags

- **database-description**—Database description packets, which are used in synchronizing the OSPF and OSPFv3 topological database.
- **error**—OSPF and OSPFv3 error packets.
- **event**—OSPF and OSPFv3 state transitions.
- **flooding**—Link-state flooding packets.
- **graceful-restart**—Graceful-restart events.
- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable.
- **ldp-synchronization**—Synchronization events between OSPF and LDP.
- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database.
- **lsa-analysis**—Link-state analysis. Specific to the Juniper Networks implementation of OSPF, Junos OS performs LSA analysis before running the shortest-path-first (SPF) algorithm. LSA analysis helps to speed the calculations performed by the SPF algorithm.
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database.
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **on-demand**—Trace demand circuit extensions.
- **packet-dump**—Content of selected packet types.
- **packets**—All OSPF packets.
- **restart-signaling**—(OSPFv2 only) Restart-signaling graceful restart events.
- **spf**—Shortest-path-first (SPF) calculations.

Global Tracing Flags

- **all**—All tracing operations.
- **general**—A combination of the **normal** and **route** trace operations.
- **normal**—All normal operations. If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions.
- **route**—Routing table changes.
- **state**—State transitions.
- **task**—Routing protocol task processing.
- **timer**—Routing protocol timer processing.

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

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

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-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.

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

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

Default: 128 KB

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

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

traceoptions (Protocols RIP)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rip], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip], [edit protocols rip], [edit routing-instances <i>routing-instance-name</i> protocols rip]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Set RIP protocol-level tracing options.



NOTE: The **traceoptions** statement is not supported on QFabric systems.

Default The default RIP protocol-level trace options are inherited from the global **traceoptions** statement.

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

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIP tracing output in the file `/var/log/rip-log`.

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

Range: 2 through 1000 files

Default: 10 files

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

RIP Tracing Options

- **auth**—RIP authentication
- **error**—RIP error packets

- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets such as request, poll, and poll entry packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

Global Tracing Options

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

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

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

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

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-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 must also specify a maximum number of trace files with the **files** option.

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

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

Default: 128 KB

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

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

traceoptions (Protocols RIPng)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ripng],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng],</p> <p>[edit protocols ripng],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>
Description	Set RIPng protocol-level tracing options.
Default	The default RIPng protocol-level trace options are inherited from the global traceoptions statement.
Options	<p>disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIPng tracing output in the file <code>/var/log/ripng-log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 10 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <p>RIPng Tracing Options</p> <ul style="list-style-type: none"> • error—RIPng error packets • expiration—RIPng route expiration processing • holddown—RIPng hold-down processing • nsr-synchronization—Nonstop routing synchronization events • packets—All RIPng packets

- **request**—RIPng information packets such as request, poll, and poll entry packets
- **trigger**—RIPng triggered updates
- **update**—RIPng update packets

Global Tracing Options

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

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

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

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

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

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

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-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 must also specify a maximum number of trace files with the **files** option.

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

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

Default: 128 KB

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

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

- Related Documentation**
- [Example: Tracing RIPv2 Protocol Traffic](#)

traceoptions (Routing Options)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <disable>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p>nsr-synchronization and nsr-packet flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>nsr-synchronization flag for RIP and RIPv6 added in Junos OS Release 9.0.</p> <p>nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p>nsr-synchronization flag for PIM added in Junos OS Release 9.3.</p> <p>nsr-synchronization flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>nsr-synchronization flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>Values:</p> <p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and</p>

so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization-detail**—(MPLS only) Tracing operations for nonstop active routing in detail
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-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. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

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

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

Default: 128 KB

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

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

Related Documentation	<ul style="list-style-type: none">• Example: Tracing Global Routing Protocol Operations• Tracing Nonstop Active Routing Synchronization Events
------------------------------	---

traffic-engineering (OSPF)

Syntax	<pre> traffic-engineering { <advertise-unnumbered-interfaces>; <credibility-protocol-preference>; ignore-lsp-metrics; multicast-rpf-routes; no-topology; shortcuts { lsp-metric-into-summary; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>multicast-rpf-routes option introduced in Junos OS Release 7.5.</p> <p>advertise-unnumbered-interfaces option introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4.</p> <p>Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>credibility-protocol-preference statement introduced in Junos OS Release 9.4.</p> <p>credibility-protocol-preference statement introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Enable the OSPF traffic engineering features.
Default	Traffic engineering support is disabled.
Options	<p>advertise-unnumbered-interfaces—(Optional) (OSPFv2 only) Include the link-local identifier in the link-local traffic-engineering link-state advertisement. This statement must be included on both ends of an unnumbered link to allow an ingress LER to update the link in its traffic engineering database and use it for CSPF calculations. The link-local identifier is then used by RSVP to signal unnumbered interfaces as defined in RFC 3477.</p> <p>credibility-protocol-preference—(Optional) (OSPFv2 only) Use the configured preference value for OSPF routes to calculate the traffic engineering database credibility value used to select IGP routes. Use this statement to override the default behavior, in which the traffic engineering database prefers IS-IS routes even if OSPF routes are configured with a lower, that is, preferred, preference value. For example, OSPF routes have a default preference value of 10, whereas IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502,</p>

whereas IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.

multicast-rpf-routes—(Optional) (OSPFv2 only) Install routes for multicast RPF checks into the **inet.2** routing table. The **inet.2** routing table consists of unicast routes used for multicast RPF lookup. RPF is an antispoofing mechanism used to check whether the packet is coming in on an interface that is also sending data back to the packet source.



NOTE: You must enable OSPF traffic engineering shortcuts to use the **multicast-rpf-routes** statement. You must not allow LSP advertisements into OSPF when configuring the **multicast-rpf-routes** statement.

no-topology—(Optional) (OSPFv2 only) Disable the dissemination of the link-state topology information.

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Enabling OSPF Traffic Engineering Support

transit-delay (OSPF)

Syntax	<code>transit-delay seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link],</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) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Set the estimated time required to transmit a link-state update on the interface. When calculating this time, make sure to account for transmission and propagation delays.</p> <p>You should never have to modify the transit delay time.</p>
Options	<p>seconds—Estimated time, in seconds.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 1 second</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring OSPF Timers • Configuring RSVP and OSPF for LMP Peer Interfaces

type (Protocols BGP)

Syntax	<code>type type;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Specify the type of BGP peer group.</p> <p>When configuring a BGP group, you can indicate whether the group is an IBGP group or an EBGP group. All peers in an IBGP group are in the same AS, while peers in an EBGP group are in different ASs and normally share a subnet.</p>
Options	<p>type—Type of group:</p> <ul style="list-style-type: none">• external—External group, which allows inter-AS BGP routing• internal—Internal group, which allows intra-AS BGP routing
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• BGP Configuration Guide

type-7

Syntax	type-7;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Flood Type 7 default link-state advertisements (LSAs) if the no-summaries statement is configured.</p> <p>By default, when the no-summaries statement is configured, a Type 3 LSA is injected into not-so-stubby areas (NSSAs) for Junos OS Release 5.0 and later. To support backward compatibility with earlier Junos OS releases, include the type-7 statement. This statement enables NSSA ABRs to advertise a Type 7 default LSA into the NSSA if you have also included the no-summaries statement in the configuration.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Areas and Router Functionality Overview • Example: Configuring OSPF Not-So-Stubby Areas • no-summaries on page 384

update-interval (Protocols RIP)

Syntax	<code>update-interval seconds;</code>
Hierarchy Level	<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>],</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>],</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>],</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>],</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>
Release Information	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the interval at which routes learned by RIP are sent to neighbors. This timer controls the interval between routing updates. This timer is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can happen if all routing devices update their neighbors simultaneously.
Options	seconds —Estimated time to wait before making updates to the routing table. Range: 10 through 60 seconds Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring RIP Timers

update-interval (Protocols RIPng)

Syntax	update-interval <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ripng], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng], [edit protocols ripng], [edit routing-instances <i>routing-instance-name</i> protocols ripng]
Release Information	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for routing instances introduced in Junos OS Release 9.0.
Description	Configure the interval at which routes learned by RIPng are sent to neighbors.
Options	seconds —Estimated time to wait before making updates to the routing table. Range: 10 through 60 seconds Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Configuring RIP Timers

upstream-interface

Syntax	<code>upstream-interface [<i>interface-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-igmp-proxy],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-mld-proxy],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options multicast pim-to-igmp-proxy],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options multicast pim-to-mld-proxy],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-igmp-proxy],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast pim-to-mld-proxy],</code> <code>[edit routing-options multicast pim-to-igmp-proxy],</code> <code>[edit routing-options multicast pim-to-mld-proxy]</code>
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 9.6 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.
Description	Configure at least one, but not more than two, upstream interfaces on the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain. The RP routing device translates PIM join or prune messages into corresponding IGMP report or leave messages (if you include the pim-to-igmp-proxy statement), or into corresponding MLD report or leave messages (if you include the pim-to-mld-proxy statement). The routing device then proxies the IGMP or MLD report or leave messages to one or both upstream interfaces to forward IPv4 multicast traffic (for IGMP) or IPv6 multicast traffic (for MLD) across the PIM domains.
Options	<i>interface-names</i> —Names of one or two upstream interfaces to which the RP routing device proxies IGMP or MLD report or leave messages for transmission of multicast traffic across PIM domains. You can specify a maximum of two upstream interfaces on the RP routing device. To configure a set of two upstream interfaces, specify the full interface names, including all physical and logical address components, within square brackets (<code>[]</code>).
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring PIM-to-IGMP Message TranslationConfiguring PIM-to-MLD Message Translation

virtual-link

Syntax	<pre>virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i> { disable; authentication key <key-id identifier>; dead-interval <i>seconds</i>; hello-interval <i>seconds</i>; ipsec-sa <i>name</i>; retransmit-interval <i>seconds</i>; transit-delay <i>seconds</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) <i>area area-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <i>area area-id</i>], [edit protocols (ospf ospf3) <i>area area-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf <i>area area-id</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>For backbone areas only, create a virtual link to use in place of an actual physical link. All area border routers and other routing devices on the backbone must be contiguous. If this is not possible and there is a break in OSPF connectivity, use virtual links to create connectivity to the OSPF backbone. When configuring virtual links, you must configure links on the two routing devices that form the end points of the link, and both of these routing devices must be area border routers. You cannot configure links through stub areas.</p>
Options	<p>neighbor-id <i>router-id</i>—IP address of the routing device at the remote end of the virtual link.</p> <p>transit-area <i>area-id</i>—Area identifier of the area through which the virtual link transits. Virtual links are not allowed to transit the backbone area.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Areas and Router Functionality Overview • Example: Configuring OSPF Virtual Links

wide-metrics-only

Syntax	wide-metrics-only;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>], [edit protocols isis level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure IS-IS to generate metric values greater than 63 on a per IS-IS level basis.</p> <p>Normally, IS-IS metrics can have values up to 63, and IS-IS generates two type, length, and value (TLV) tuples, one for an IS-IS adjacency and the second for an IP prefix. To allow IS-IS to support traffic engineering, a second pair of TLVs has been added to IS-IS, one for IP prefixes and the second for IS-IS adjacency and traffic engineering information. With these TLVs, IS-IS metrics can have values up to 16,777,215 ($2^{24} - 1$).</p> <p>To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, include the wide-metrics-only statement.</p>
Default	By default, Junos OS supports the sending and receiving of wide metrics. Junos OS allows a maximum metric value of 63 and generates both pairs of TLVs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Enabling Wide IS-IS Metrics for Traffic Engineeringte-metric

CHAPTER 5

Operational Commands for Layer 3 Protocols

clear (ospf | ospf3) database

Syntax clear (ospf | ospf3) database
<advertising-router (*router-id* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<logical-system (all | *logical-system-name*)>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<purge>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>

Syntax (EX Series Switch and QFX Series) clear (ospf | ospf3) database
<advertising-router (*router-id* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<purge>
<router>

Release Information Command introduced before Junos OS Release 7.4.
advertising-router *router-id*, **area** *area-id*, **asbrsummary**, **external**, **inter-area-prefix**, **inter-area-router**, **intra-area-prefix**, **link-local**, **lsa-id** *lsa-id*, **netsummary**, **network**, **nssa**, **opaque-area**, and **router** options added in Junos OS Release 8.3. You must use the **purge** command with these options.
Command introduced in Junos OS Release 9.0 for EX Series switches.
realm option added in Junos OS Release 9.2.
advertising-router (*router-id* | **self**) option added in Junos OS Release 9.5.
advertising-router (*router-id* | **self**) option introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description With the master Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and sync the new database with the master Routing Engine. You can also use the **purge** command with any of the options to discard rather than delete the specified LSA entries.



CAUTION: This command is useful only for testing. Use it with care, because it causes significant network disruption.

Options **none**—Delete all LSAs other than the system's own LSAs, which are regenerated. To resynchronize the database, the system destroys all adjacent neighbors that are in the state **EXSTART** or higher. The neighbors are then reacquired and the databases are synchronized.

advertising-router (*router-id* | **self**)—(Optional) Discard entries for the LSA entries advertised by the specified routing device or by this routing device.

area *area-id*—(Optional) Discard entries for the LSAs in the specified area.

asbrsummary—(Optional) Discard summary AS boundary router LSA entries.

external—(Optional) Discard external LSAs.

instance *instance-name*—(Optional) Delete or discard entries for the specified routing instance only.

inter-area-prefix—(OSPFv3 only) (Optional) Discard interarea prefix LSAs.

inter-area-router—(OSPFv3 only) (Optional) Discard interarea router LSAs.

intra-area-prefix—(OSPFv3 only) (Optional) Discard intra-area prefix LSAs.

logical-system (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

link-local—(Optional) Delete link-local LSAs.

lsa-id *lsa-id*—(Optional) Discard the LSA entries with the specified LSA identifier.

netsummary—(Optional) Discard summary network LSAs.

network—(Optional) Discard network LSAs.

nssa—(Optional) Discard not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Discard opaque area-scope LSAs.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Delete the entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Discard router LSAs.

purge—(Optional) Discard all entries in the link-state advertisement database. All link-state advertisements are set to **MAXAGE** and are flooded. The database is repopulated when the originators of the link-state advertisements receive the **MAXAGE** link-state advertisements and reissue them.

Required Privilege Level

clear

Related Documentation

- [show ospf database on page 569](#)
- [show ospf3 database on page 558](#)

List of Sample Output [clear ospf database on page 422](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf database

```
user@host> clear ospf database
```

clear (ospf | ospf3) io-statistics

Syntax	clear (ospf ospf3) io-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	clear (ospf ospf3) io-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Open Shortest Path First (OSPF) input and output statistics.
Options	none —Clear OSPF input and output statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
List of Sample Output	clear ospf io-statistics on page 423
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf io-statistics

```
user@host> clear ospf io-statistics
```

clear (ospf | ospf3) neighbor

Syntax	clear (ospf ospf3) neighbor <area <i>area-id</i> > <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> <neighbor> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	clear (ospf ospf3) neighbor <area <i>area-id</i> > <instance <i>instance-name</i> > <interface <i>interface-name</i> > <neighbor>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Tear down Open Shortest Path First (OSPF) neighbor connections.
Options	none —Tear down OSPF connections with all neighbors for all routing instances. area <i>area-id</i> —(Optional) Tear down neighbor connections for the specified area only. instance <i>instance-name</i> —(Optional) Tear down neighbor connections for the specified routing instance only. interface <i>interface-name</i> —(Optional) Tear down neighbor connections for the specified interface only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. neighbor —(Optional) Clear the state of the specified neighbor only. realm (ipv4-multicast ipv4-unicast ipv6-multicast) —(Optional) (OSPFv3 only) Clear the state of the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show (ospf ospf3) neighbor on page 458
List of Sample Output	clear ospf neighbor on page 425
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf neighbor

```
user@host> clear ospf neighbor
```

clear (ospf | ospf3) statistics

Syntax	clear (ospf ospf3) statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	clear (ospf ospf3) statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Open Shortest Path First (OSPF) statistics.
Options	none —Clear OSPF statistics. instance <i>instance-name</i> —(Optional) Clear statistics for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. realm (ipv4-multicast ipv4-unicast ipv6-multicast) —(Optional) (OSPFv3 only) Clear statistics for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">show (ospf ospf3) statistics on page 475
List of Sample Output	clear ospf statistics on page 426
Output Fields	See show (ospf ospf3) statistics for an explanation of output fields.

Sample Output

clear ospf statistics

The following sample output displays OSPF statistics before and after the **clear ospf statistics** command is entered:

```
user@host> show ospf statistics
```

Packet type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	3254	2268	3	1
DbD	41	46	0	0


```

LSReq          8          7          0          0
LSUpdate       212        154         0          0
LSAck          65         98          0          0

DBDs retransmitted :          3, last 5 seconds :      0
LSAs flooded       :         12, last 5 seconds :      0
LSAs flooded high-prio :      0, last 5 seconds :      0
LSAs retransmitted :          0, last 5 seconds :      0
LSAs transmitted to nbr:      3, last 5 seconds :      0
LSAs requested     :          5, last 5 seconds :      0
LSAs acknowledged :         19, last 5 seconds :      0

Flood queue depth :          0
Total rexmit entries :        0
db summaries      :          0
lsreq entries     :          0

Receive errors:
  626 subnet mismatches

```

```
user@host> clear ospf statistics
```

```

user@host> show ospf statistics
Packet type      Total
                Sent  Received
Hello            3      1
  DbD            0      0
  LSReq          0      0
LSUpdate         0      0
LSAck            0      0

                Last 5 seconds
                Sent  Received
Hello            3      1
  DbD            0      0
  LSReq          0      0
LSUpdate         0      0
LSAck            0      0

DBDs retransmitted :          0, last 5 seconds :      0
LSAs flooded       :          0, last 5 seconds :      0
LSAs flooded high-prio :      0, last 5 seconds :      0
LSAs retransmitted :          0, last 5 seconds :      0
LSAs transmitted to nbr:      0, last 5 seconds :      0
LSAs requested     :          0, last 5 seconds :      0
LSAs acknowledged :          0, last 5 seconds :      0

Flood queue depth :          0
Total rexmit entries :        0
db summaries      :          0
lsreq entries     :          0

Receive errors:
  None

```

clear bgp damping

Syntax	clear bgp damping <logical-system (all <i>logical-system-name</i>)> < <i>prefix</i> >
Syntax (EX Series Switch and QFX Series)	clear bgp damping < <i>prefix</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear BGP route flap damping information.
Options	none —Clear all BGP route flap damping information. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>prefix</i> —(Optional) Clear route flap damping information for only the specified destination prefix.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show policy damping on page 577• show route damping on page 614
List of Sample Output	clear bgp damping on page 428
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bgp damping

```
user@host> clear bgp damping
```

clear bgp neighbor

Syntax	<pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre>
Syntax (EX Series Switch and QFX Series)	<pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Perform one of the following tasks:</p> <ul style="list-style-type: none"> • Change the state of one or more BGP neighbors to IDLE. For neighbors in the ESTABLISHED state, this command drops the TCP connection to the neighbors and then reestablishes the connection. • (soft or soft-inbound keyword only) Reapply export policies or import policies, respectively, and send refresh updates to one or more BGP neighbors without changing their state.
Options	<p>none—Change the state of all BGP neighbors to IDLE.</p> <p>as <i>as-number</i>—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).</p> <p>instance <i>instance-name</i>—(Optional) Apply this command only to neighbors for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.</p> <p>soft—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.</p> <p>soft-inbound—(Optional) Reapply any import policies and send refresh updates to neighbors without clearing the state.</p>

soft-minimum-igp—(Optional) Provides soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.

Required Privilege Level clear

Related Documentation • [show bgp neighbor on page 496](#)

List of Sample Output [clear bgp neighbor on page 430](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear bgp neighbor](#)

```
user@host> clear bgp neighbor
```

clear bgp table

Syntax	<code>clear bgp table <i>table-name</i></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switch and QFX Series)	<code>clear bgp table <i>table-name</i></code>
Release Information	Command introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Request that BGP refresh routes in a specified routing table.
Options	<code>logical-system (all <i>logical-system-name</i>)</code> —(Optional) Perform this operation on all logical systems or on a particular logical system. <code>table-name</code> —Request that BGP refresh routes in the specified table.
Additional Information	In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the clear bgp table command to request that BGP refresh routes in a VPN instance table.
Required Privilege Level	clear
List of Sample Output	clear bgp table private.inet.0 on page 431 clear bgp table inet.6 logical-system all on page 431 clear bgp table private.inet.6 logical-system ls1 on page 431 clear bgp table logical-system all inet.0 on page 431 clear bgp table logical-system ls2 private.inet.0 on page 432
Output Fields	This command produces no output.

Sample Output

[clear bgp table private.inet.0](#)

```
user@host> clear bgp table private.inet.0
```

[clear bgp table inet.6 logical-system all](#)

```
user@host> clear bgp table inet.6 logical-system all
```

[clear bgp table private.inet.6 logical-system ls1](#)

```
user@host> clear bgp table private.inet.6 logical-system ls1
```

[clear bgp table logical-system all inet.0](#)

```
user@host> clear bgp table logical-system all inet.0
```

clear bgp table logical-system ls2 private.inet.0

```
user@host> clear bgp table logical-system ls2 private.inet.0
```

clear ipv6 neighbors

Syntax	clear ipv6 neighbors <all host <i>hostname</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Clear IPv6 neighbor cache information.
Options	none —Clear all IPv6 neighbor cache information. all —(Optional) Clear all IPv6 neighbor cache information. host <i>hostname</i> —(Optional) Clear the information for the specified IPv6 neighbors.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show ipv6 neighbors on page 515
List of Sample Output	clear ipv6 neighbors on page 433
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
```

clear isis adjacency

Syntax	clear isis adjacency <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> <neighbor>
Syntax (EX Series Switches and QFX Series)	clear isis adjacency <instance <i>instance-name</i> > <interface <i>interface-name</i> > <neighbor>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Remove entries from the IS-IS adjacency database.
Options	none —Remove all entries from the adjacency database. instance <i>instance-name</i> —(Optional) Clear all adjacencies for the specified routing instance only. interface <i>interface-name</i> —(Optional) Clear all adjacencies for the specified interface only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. neighbor —(Optional) Clear adjacencies for the specified neighbor only.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis adjacency on page 517
List of Sample Output	clear isis adjacency on page 434
Output Fields	See show isis adjacency for an explanation of output fields.

Sample Output

clear isis adjacency

The following sample output displays IS-IS adjacency database information before and after the **clear isis adjacency** command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State      Hold (secs) SNPA
so-1/0/0.0     karaku1         3 Up         26
so-1/1/3.0     1921.6800.5080 3 Up         23
```



```
so-5/0/0.0    1921.6800.5080 3 Up                19
```


```
user@host> clear isis adjacency karakul
```

```
user@host> show isis adjacency
```

```
IS-IS adjacency database:
```

Interface	System	L State	Hold (secs)	SNPA
so-1/0/0.0	karakul	3 Initializing	26	
so-1/1/3.0	1921.6800.5080	3 Up	24	
so-5/0/0.0	1921.6800.5080	3 Up	21	

clear isis database

Syntax	clear isis database <entries> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <purge>
Syntax (EX Series Switches and QFX Series)	clear isis database <entries> <instance <i>instance-name</i> > <purge>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. purge option introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Remove the entries from the IS-IS link-state database, which contains prefixes and topology information. You can also use purge with any of the options to initiate a network-wide purge of link-state PDUs rather than the local deletion of entries from the IS-IS link-state database.
<div><div>CAUTION: In a production network, the purge command option might cause short-term network-wide traffic disruptions.</div></div>	
Options	none —Remove all entries from the IS-IS link-state database for all routing instances. entries —(Optional) Name of the database entry. instance <i>instance-name</i> —(Optional) Clear all entries for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. purge —(Optional) Discard all entries in the IS-IS link-state database.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis database on page 531
List of Sample Output	clear isis database on page 437
Output Fields	See show isis database for an explanation of output fields.

Sample Output

clear isis database

The following sample output displays IS-IS link-state database information before and after the **clear isis database** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID          Sequence Checksum Lifetime (secs)
crater.00-00     0x12    0x84dd             1139
1 LSPs
IS-IS level 2 link-state database:
LSP ID          Sequence Checksum Lifetime (secs)
crater.00-00     0x19    0xe92c             1134
badlands.00-00   0x16    0x1454             985
carlsbad.00-00   0x33    0x220b            1015
ranier.00-00     0x2e    0xfc31            1007
1921.6800.5066.00-00 0x11    0x7313             566
1921.6800.5067.00-00 0x14    0xd9d4             939
6 LSPs
```

```
user@host> clear isis database
```

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID          Sequence Checksum Lifetime (secs)

IS-IS level 2 link-state database:
LSP ID          Sequence Checksum Lifetime (secs)
```

clear isis overload

Syntax	clear isis overload <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	clear isis overload <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Reset the IS-IS dynamic overload bit. This command can appear to not work, continuing to display overload after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.</p> <p>When other routers detect that the overload bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes.</p>
Options	<p>none—Reset the IS-IS dynamic overload bit.</p> <p>instance <i>instance-name</i>—(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis database on page 531
List of Sample Output	clear isis overload on page 438
Output Fields	See show isis database for an explanation of output fields.

Sample Output

clear isis overload

The following sample output displays IS-IS database information before and after the **clear isis overload** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
pro3-c.00-00          0x4    0x10db    1185 L1 L2 Overload

  1 LSPs
```

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0x5	0x429f	1185	L1 L2 Overload

pro2-a.00-00	0x91e	0x2589	874	L1 L2
pro2-a.02-00	0x1	0xcbc	874	L1 L2
3 LSPs				

user@host> clear isis overload

user@host> show isis database

IS-IS level 1 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0xa	0x429e	1183	L1 L2
1 LSPs				

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0xc	0x9c39	1183	L1 L2
pro2-a.00-00	0x91e	0x2589	783	L1 L2
pro2-a.02-00	0x1	0xcbc	783	L1 L2
3 LSPs				

clear isis statistics

Syntax	clear isis statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	clear isis statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Set statistics about IS-IS traffic to zero.
Options	none —Set IS-IS traffic statistics to zero for all routing instances. instance <i>instance-name</i> —(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show isis statistics on page 555
List of Sample Output	clear isis statistics on page 440
Output Fields	See show isis statistics for an explanation of output fields.

Sample Output

clear isis statistics

The following sample output displays IS-IS statistics before and after the **clear isis statistics** command is entered:

```
user@host> show isis statistics
IS-IS statistics for merino:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	12793	12793	0	8666	719
IIH	116751	116751	0	118834	0
CSNP	203956	203956	0	204080	0
PSNP	7356	7350	6	8635	0
Unknown	0	0	0	0	0
Totals	340856	340850	6	340215	719

Total packets received: 340856 Sent: 340934

SNP queue length: 0 Drops: 0

LSP queue length: 0 Drops: 0

SPF runs: 1064
Fragments rebuilt: 1087
LSP regenerations: 436
Purges initiated: 0

user@host> clear isis statistics

user@host> show isis statistics
IS-IS statistics for merino:

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	0	0	0	0	0
IIH	3	3	0	3	0
CSNP	2	2	0	4	0
PSNP	0	0	0	0	0
Unknown	0	0	0	0	0
Totals	5	5	0	7	0

Total packets received: 5 Sent: 7

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0

SPF runs: 0
Fragments rebuilt: 0
LSP regenerations: 0
Purges initiated: 0

clear (ospf | ospf3) overload

Syntax	clear (ospf ospf3) overload <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	clear (ospf ospf3) overload <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear the Open Shortest Path First (OSPF) overload bit and rebuild link-state advertisements (LSAs).
Options	none —Clear the overload bit and rebuild LSAs for all routing instances. instance <i>instance-name</i> —(Optional) Clear the overload bit and rebuild LSAs for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
List of Sample Output	clear ospf overload on page 442
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf overload

```
user@host> clear ospf overload
```


clear rip general-statistics

Syntax	clear rip general-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	clear rip general-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear RIP general statistics.
Options	<p>none—Clear RIP general statistics.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show rip general-statistics on page 579
List of Sample Output	clear rip general-statistics on page 443
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rip general-statistics

```
user@host> clear rip general-statistics
```

clear rip statistics

Syntax	clear rip statistics <instance (all <i>instance-name</i>)> <logical-system (all <i>logical-system-name</i>)> <neighbor> <peer (all <i>address</i>)>
Syntax (EX Series Switches and QFX Series)	clear rip statistics <instance (all <i>instance-name</i>)> <neighbor>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear RIP statistics.
Options	none —Reset RIP counters for all neighbors for all routing instances. instance (all <i>instance-name</i>) —(Optional) Clear RIP statistics for all instances or for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. neighbor —(Optional) Clear RIP statistics for the specified neighbor only. peer (all <i>address</i>) —(Optional) Clear RIP statistics for a single peer or all peers.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show rip statistics on page 583
List of Sample Output	clear rip statistics on page 444
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rip statistics

```
user@host> clear rip statistics
```

clear ripng general-statistics

Syntax	clear ripng general-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	clear ripng general-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Clear RIP next generation (RIPng) general statistics.
Options	none —Clear RIPng general statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show ripng general-statistics on page 586
List of Sample Output	clear ripng general-statistics on page 445
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ripng general-statistics

```
user@host> clear ripng general-statistics
```

clear ripng statistics

Syntax	clear ripng statistics < <i>instance</i> <i>name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch)	clear ripng statistics < <i>instance</i> <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Clear RIP next-generation (RIPng) statistics.
Options	none —Reset RIPng counters for all neighbors for all routing instances. <i>instance</i> —(Optional) Reset RIPng counters for the specified instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>name</i> —(Optional) Reset RIPng counters for the specified neighbor.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show ripng statistics on page 589
List of Sample Output	clear ripng statistics on page 446
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ripng statistics

```
user@host> clear ripng statistics
```

show (ospf | ospf3) interface

Syntax	<pre>show (ospf ospf3) interface <brief detail extensive> <area <i>area-id</i>> <<i>interface-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show (ospf ospf3) interface <brief detail extensive> <area <i>area-id</i>> <<i>interface-name</i>> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>area option introduced in Junos OS Release 9.2.</p> <p>area option introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the status of OSPF interfaces.
Options	<p>none—Display standard information about the status of all OSPF interfaces for all routing instances</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display information about the interfaces that belong to the specified area.</p> <p><i>interface-name</i>—(Optional) Display information for the specified interface.</p> <p>instance <i>instance-name</i>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(OSPFv3 only) (Optional) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	<p>show ospf interface brief on page 450</p> <p>show ospf interface detail on page 450</p> <p>show ospf3 interface detail on page 450</p>

[show ospf interface detail \(When Multiarea Adjacency Is Configured\) on page 450](#)
[show ospf interface area area-id on page 451](#)
[show ospf interface extensive \(When Flooding Reduction Is Enabled\) on page 452](#)
[show ospf interface extensive \(When LDP Synchronization Is Configured\) on page 452](#)

Output Fields Table 23 on page 448 lists the output fields for the **show (ospf | ospf3) interface** command. Output fields are listed in the approximate order in which they appear.

Table 23: show (ospf | ospf3) interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface running OSPF version 2 or OSPF version 3.	All levels
State	State of the interface: BDR , Down , DR , DRother , Loop , PtToPt , or Waiting .	All levels
Area	Number of the area that the interface is in.	All levels
DR ID	Address of the area's designated router.	All levels
BDR ID	Backup designated router for a particular subnet.	All levels
Nbrs	Number of neighbors on this interface.	All levels
Type	Type of interface: LAN , NBMA , P2MP , P2P , or Virtual .	detail extensive
Address	IP address of the neighbor.	detail extensive
Mask	Netmask of the neighbor.	detail extensive
Prefix-length	(OSPFv3) IPv6 prefix length, in bits.	detail extensive
OSPF3-Intf-Index	(OSPFv3) OSPF version 3 interface index.	detail extensive
MTU	Interface maximum transmission unit (MTU).	detail extensive
Cost	Interface cost (metric).	detail extensive
DR addr	Address of the designated router.	detail extensive
BDR addr	Address of the backup designated router.	detail extensive
Adj count	Number of adjacent neighbors.	detail extensive
Secondary	Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface for only one area.	detail extensive
Flood Reduction	Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the DoNotAge bit set. As a result, LSAs are refreshed only when a change occurs.	extensive

Table 23: show (ospf | ospf3) interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Priority	Router priority used in designated router (DR) election on this interface.	detail extensive
Flood list	List of link-state advertisements (LSAs) that might be about to flood this interface.	extensive
Ack list	Acknowledgment list. List of pending acknowledgments on this interface.	extensive
Descriptor list	List of packet descriptors.	extensive
Hello	Configured value for the hello timer.	detail extensive
Dead	Configured value for the dead timer.	detail extensive
Auth type	(OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: <ul style="list-style-type: none"> • MD5—The MD5 mechanism is configured in accordance with RFC 2328. • None—No authentication method is configured. • Password—A simple password (RFC 2328) is configured. 	detail extensive
Topology	(Multiarea adjacency) Name of topology: default or <i>name</i> .	
LDP sync state	(OSPFv2 and LDP synchronization) Current state of LDP synchronization: in sync , in holddown , and not supported .	extensive
reason	(OSPFv2 and LDP synchronization) Reason for the current state of LDP synchronization. The LDP session might be up or down, or adjacency might be up or down.	extensive
config holdtime	(OSPFv2 and LDP synchronization) Configured value of the hold timer. If the state is not synchronized, and the hold time is not infinity, the remaining field displays the number of seconds that remain until the configured hold timer expires.	extensive
IPSec SA name	(OSPFv2) Name of the IPSec security association name.	detail extensive
Active key ID	(OSPFv2 and MD5) Number from 0 to 255 that uniquely identifies an MD5 key.	detail extensive
Start time	(OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as Start time 1970 Jan 01 00:00:00 PST .	detail extensive
ReXmit	Configured value for the Retransmit timer.	detail extensive
Stub, Not Stub, or Stub NSSA	Type of area.	detail extensive

Sample Output

show ospf interface brief

```
user@host> show ospf interface brief
```

Intf	State	Area	DR ID	BDR ID	Nbrs
at-5/1/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
ge-2/3/0.0	DR	0.0.0.0	192.168.4.16	192.168.4.15	1
lo0.0	DR	0.0.0.0	192.168.4.16	0.0.0.0	0
so-0/0/0.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-6/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-6/0/2.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-6/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

show ospf interface detail

```
user@host> show ospf interface detail
```

Interface	State	Area	DR ID	BDR ID	Nbrs
fe-0/0/1.0	BDR	0.0.0.0	192.168.37.12	10.255.245.215	1

Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40
 DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128
 Hello 10, Dead 40, ReXmit 5, Not Stub

Interface	State	Area	DR ID	BDR ID	Nbrs
tl-0/2/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	0

Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604
 Adj count 0
 Hello 10, Dead 40, ReXmit 5, Not Stub
 Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST
 IPsec SA Name: sa

show ospf3 interface detail

```
user@host> show ospf3 interface so-0/0/3.0 detail
```

Interface	State	Area	DR-ID	BDR-ID	Nbrs
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64
 OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1
 Hello 10, Dead 40, ReXmit 5, Not Stub

show ospf interface detail (When Multiarea Adjacency Is Configured)

```
user@host> show ospf interface detail
```

```
regress@router> show ospf interface detail
```

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.255.245.2	0.0.0.0	0

Type: LAN, Address: 127.0.0.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
 DR addr: 127.0.0.1, Adj count: 0, Priority: 128
 Hello: 10, Dead: 40, ReXmit: 5, Not Stub
 Auth type: None
 Topology default (ID 0) -> Cost: 0

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.255.245.2	0.0.0.0	0

Type: LAN, Address: 10.255.245.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0
 DR addr: 10.255.245.2, Adj count: 0, Priority: 128
 Hello: 10, Dead: 40, ReXmit: 5, Not Stub
 Auth type: None
 Topology default (ID 0) -> Cost: 0

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1


```

Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      0

Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      0

Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0      PtToPt  1.1.1.1      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  1.1.1.1      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  2.2.2.2      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  2.2.2.2      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

```

show ospf interface area area-id

```
user@host> show ospf interface area 1.1.1.1
```

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/0/0.0	PtToPt	1.1.1.1	0.0.0.0	0.0.0.0	1
so-1/0/0.0	PtToPt	1.1.1.1	0.0.0.0	0.0.0.0	1

show ospf interface extensive
(When Flooding Reduction Is Enabled)

```
user@host> show ospf interface extensive
Interface      State  Area      DR ID      BDR ID      Nbrs
fe-0/0/0.0     PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     0

Type: P2P, Address: 10.10.10.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
Adj count: 0
Secondary, Flood Reduction
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
```

show ospf interface extensive
(When LDP Synchronization Is Configured)

```
user@host> show ospf interface extensive
Interface      State  Area      DR ID      BDR ID
Nbrs
so-1/0/3.0     Down   0.0.0.0    0.0.0.0    0.0.0.0
0

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 65535
Adj count: 0
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
LDP sync state: in holddown, for: 00:00:08, reason: LDP down during config
config holddtime: 10 seconds, remaining: 1
```

show (ospf | ospf3) io-statistics

Syntax	show (ospf ospf3) io-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) io-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display Open Shortest Path First (OSPF) input and output statistics.
Options	none —Display OSPF input and output statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear (ospf ospf3) statistics on page 426
List of Sample Output	show ospf io-statistics on page 453
Output Fields	Table 24 on page 453 lists the output fields for the show ospf io-statistics command. Output fields are listed in the approximate order in which they appear.

Table 24: show (ospf | ospf3) io-statistics Output Fields

Field Name	Field Description
Packets read	Number of OSPF packets read since the last time the routing protocol was started.
average per run	Total number of packets divided by the total number of times the OSPF read operation is scheduled to run.
max run	Maximum number of packets for a given run among all scheduled runs.
Receive errors	Number of faulty packets received with errors.

Sample Output

show ospf io-statistics

```
user@host> show ospf io-statistics
```

```
Packets read: 7361, average per run: 1.00, max run: 1
```

Receive errors:
None

show (ospf | ospf3) log

Syntax	show (ospf ospf3) log <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <topology <i>topology-name</i> >
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) log <instance <i>instance-name</i> > <topology <i>topology-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. topology option introduced in Junos OS Release 9.0. topology option introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations.
Options	<p>none—Display entries in the OSPF log of SPF calculations for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display entries for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>topology <i>topology-name</i>—(Optional) (OSPFv2 only) Display entries for the specified topology.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(OSPFv3 only) (Optional) Display entries for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	show ospf log on page 456 show ospf log topology voice on page 456
Output Fields	Table 25 on page 455 lists the output fields for the show (ospf ospf3) log command. Output fields are listed in the approximate order in which they appear.

Table 25: show (ospf | ospf3) log Output Fields

Field Name	Field Description
When	Time, in weeks (w) and days (d), since the SPF calculation was made.

Table 25: show (ospf | ospf3) log Output Fields (*continued*)

Field Name	Field Description
Type	Type of calculation: Cleanup, External, Interarea, NSSA, Redist, SPF, Stub, Total, or Virtuallink.
Elapsed	Amount of time, in seconds, that elapsed during the operation, or the time required to complete the SPF calculation. The start time is the time displayed in the When field.

Sample Output

show ospf log

```

user@host> show ospf log
When          Type          Elapsed
1w4d 17:25:58 Stub          0.000017
1w4d 17:25:58 SPF           0.000070
1w4d 17:25:58 Stub          0.000019
1w4d 17:25:58 Interarea     0.000054
1w4d 17:25:58 External      0.000005
1w4d 17:25:58 Cleanup       0.000203
1w4d 17:25:58 Total         0.000537
1w4d 17:24:48 SPF           0.000125
1w4d 17:24:48 Stub          0.000017
1w4d 17:24:48 SPF           0.000100
1w4d 17:24:48 Stub          0.000016
1w4d 17:24:48 Interarea     0.000056
1w4d 17:24:48 External      0.000005
1w4d 17:24:48 Cleanup       0.000238
1w4d 17:24:48 Total         0.000600
...

```

show ospf log topology voice

```

user@host> show ospf log topology voice
Topology voice SPF log:

    Last instance of each event type
When          Type          Elapsed
00:06:11      SPF           0.000116
00:06:11      Stub          0.000114
00:06:11      Interarea     0.000126
00:06:11      External      0.000067
00:06:11      NSSA          0.000037
00:06:11      Cleanup       0.000186

    Maximum length of each event type
When          Type          Elapsed
00:13:43      SPF           0.000140
00:13:33      Stub          0.000116
00:13:43      Interarea     0.000128
00:13:33      External      0.000075
00:13:38      NSSA          0.000039
00:13:53      Cleanup       0.000657

    Last 100 events

```

When	Type	Elapsed
00:13:53	SPF	0.000090
00:13:53	Stub	0.000041
00:13:53	Interarea	0.000123
00:13:53	External	0.000040
00:13:53	NSSA	0.000038
00:13:53	Cleanup	0.000657
00:13:53	Total	0.001252
.		
.		
00:06:11	SPF	0.000116
00:06:11	Stub	0.000114
00:06:11	Interarea	0.000126
00:06:11	External	0.000067
00:06:11	NSSA	0.000037
00:06:11	Cleanup	0.000186
00:06:11	Total	0.000818

show (ospf | ospf3) neighbor

Syntax	<pre>show (ospf ospf3) neighbor <brief detail extensive> <area <i>area-id</i>> <instance (all <i>instance-name</i>)> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show (ospf ospf3) neighbor <brief detail extensive> <area <i>area-id</i>> <instance (all <i>instance-name</i>)> <interface <i>interface-name</i>> <neighbor></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>instance all option introduced in Junos OS Release 9.1.</p> <p>instance all option introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>area, interface, and realm options introduced in Junos OS Release 9.2.</p> <p>area and interface options introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Display information about OSPF neighbors.</p> <p>CPU utilization might increase while the device learns its OSPF neighbors. We recommend that you use the show (ospf ospf3) neighbor command after the device learns and establishes OSPF neighbor adjacencies. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the show (ospf ospf3) neighbor command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the CLI.</p>
Options	<p>none—Display standard information about all OSPF neighbors for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display information about the OSPF neighbors for the specified area.</p> <p>instance (all <i>instance-name</i>)—(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display information about OSPF neighbors for the specified logical interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>

neighbor—(Optional) Display information about the specified OSPF neighbor.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Display information about the OSPF neighbors for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level view

Related Documentation

- [clear \(ospf | ospf3\) neighbor on page 424](#)

List of Sample Output

- [show ospf neighbor brief on page 461](#)
- [show ospf neighbor detail on page 461](#)
- [show ospf neighbor extensive on page 462](#)
- [show ospf3 neighbor detail on page 463](#)
- [show ospf neighbor area area-id on page 463](#)
- [show ospf neighbor interface interface-name on page 463](#)
- [show ospf3 neighbor instance all \(OSPFv3 Multiple Family Address Support Enabled\) on page 463](#)

Output Fields [Table 26 on page 459](#) lists the output fields for the **show (ospf | ospf3) neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 26: show (ospf | ospf3) neighbor Output Fields

Field Name	Field Description	Level of Output
Address	Address of the neighbor.	All levels
Interface	Interface through which the neighbor is reachable.	All levels

Table 26: show (ospf | ospf3) neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	<p>State of the neighbor:</p> <ul style="list-style-type: none"> • Attempt—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor. • Down—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the Down state, although at a reduced frequency. • Exchange—Routing device is describing its entire link-state database by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged. • ExStart—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the master, and to determine the initial sequence number. • Full—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements. • Init—A hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state might occur, for example, because the routing device itself did not appear in the neighbor's hello packet. • Loading—Link-state request packets are sent to the neighbor to acquire more recent advertisements that have been discovered (but not yet received) in the Exchange state. • 2Way—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state 2Way or greater. 	All levels
ID	Router ID of the neighbor.	All levels
Pri	Priority of the neighbor to become the designated router.	All levels
Dead	Number of seconds until the neighbor becomes unreachable.	All levels
Link state acknowledgment list	Number of link-state acknowledgments received.	extensive
Link state retransmission list	<p>Total number of link-state advertisements retransmitted. For extensive output only, the following information is also displayed:</p> <ul style="list-style-type: none"> • Type—Type of link advertisement: ASBR, Sum, Extern, Network, NSSA, OpagArea, Router, or Summary. • LSA ID—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device. • Adv rtr—Address of the routing device that sent the advertisement. • Seq—Link sequence number of the advertisement. 	detail extensive

Table 26: show (ospf | ospf3) neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Neighbor-address	(OSPFv3 only) If the neighbor uses virtual links, the Neighbor-address is the site-local, local, or global address. If the neighbor uses a physical interface, the Neighbor-address is an IPv6 link-local address.	detail extensive
area	Area that the neighbor is in.	detail extensive
OSPF3-Intf-Index	(OSPFv3 only) Displays the OSPFv3 interface index.	detail extensive
opt	Option bits received in the hello packets from the neighbor.	detail extensive
DR or DR-ID	Address of the designated router.	detail extensive
BDR or BDR-ID	Address of the backup designated router.	detail extensive
Up	Length of time since the neighbor came up.	detail extensive
adjacent	Length of time since the adjacency with the neighbor was established.	detail extensive

Sample Output

show ospf neighbor brief

```

user@host> show ospf neighbor brief
  Address      Intf      State      ID          Pri  Dead
192.168.254.225 fxp3.0    2Way       10.250.240.32 128  36
192.168.254.230 fxp3.0    Full       10.250.240.8  128  38
192.168.254.229 fxp3.0    Full       10.250.240.35 128  33
10.1.1.129      fxp2.0    Full       10.250.240.12 128  37
10.1.1.131      fxp2.0    Full       10.250.240.11 128  38
10.1.2.1        fxp1.0    Full       10.250.240.9  128  32
10.1.2.81       fxp0.0    Full       10.250.240.10 128  33

```

show ospf neighbor detail

```

user@host> show ospf neighbor detail
  Address      Interface      State      ID          Pri  Dead
10.5.1.2      ge-1/2/0.1    Full       10.5.1.2    128  37
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:28, adjacent 05:17:36
Link state acknowledgment list: 3 entries

Link state retransmission list: 9 entries

10.5.10.2      ge-1/2/0.10    ExStart    10.5.1.38   128  34
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:28
master, seq 0xac1530f8, rexmit DBD in 3 sec
rexmit LSREQ in 0 sec
10.5.11.2      ge-1/2/0.11    Full       10.5.1.42   128  38
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:28, adjacent 05:26:46
Link state retransmission list: 1 entries

```

```

10.5.12.2      ge-1/2/0.12      ExStart  10.5.1.46      128    33
area 0.0.0.1, opt 0x42, DR 10.5.12.2, BDR 10.5.12.1
Up 06:09:28
master, seq 0xac188a68, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec

```

show ospf neighbor extensive

```

user@host> show ospf neighbor extensive
Address      Interface      State      ID      Pri  Dead
10.5.1.2      ge-1/2/0.1     Full       10.5.1.2 128   33
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:42, adjacent 05:17:50
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
Summary    10.8.56.0    172.25.27.82 0x8000004d
Router     10.5.1.94    10.5.1.94    0x8000005c
Network    10.5.24.2    10.5.1.94    0x80000036
Summary    10.8.57.0    172.25.27.82 0x80000024
Extern     1.10.90.0    10.8.1.2     0x80000041
Extern     1.4.109.0    10.6.1.2     0x80000041
Router     10.5.1.190   10.5.1.190   0x8000005f
Network    10.5.48.2    10.5.1.190   0x8000003d
Summary    10.8.58.0    172.25.27.82 0x8000004d
Extern     1.10.91.0    10.8.1.2     0x80000041
Extern     1.4.110.0    10.6.1.2     0x80000041
Router     10.5.1.18    10.5.1.18    0x8000005f
Network    10.5.5.2     10.5.1.18    0x80000033
Summary    10.8.59.0    172.25.27.82 0x8000003a
Summary    10.8.62.0    172.25.27.82 0x80000025

10.5.10.2     ge-1/2/0.10     ExStart  10.5.1.38      128    38
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:42
master, seq 0xac1530f8, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec
10.5.11.2     ge-1/2/0.11     Full     10.5.1.42      128    33
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:42, adjacent 05:27:00
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
Summary    10.8.58.0    172.25.27.82 0x8000004d

```

Extern	1.10.91.0	10.8.1.2	0x80000041
Extern	1.1.247.0	10.5.1.2	0x8000003f
Extern	1.4.110.0	10.6.1.2	0x80000041
Router	10.5.1.18	10.5.1.18	0x8000005f
Network	10.5.5.2	10.5.1.18	0x80000033
Summary	10.8.59.0	172.25.27.82	0x8000003a

show ospf3 neighbor detail

```
user@host> show ospf3 neighbor detail
ID          Interface          State    Pri    Dead
10.255.71.13 fe-0/0/2.0          Full     128    30
Neighbor-address fe80::290:69ff:fe9b:e002
Area 0.0.0.0, opt 0x13, OSPF3-Intf-Index 2
DR-ID 10.255.71.13, BDR-ID 10.255.71.12
Up 02:51:43, adjacent 02:51:43
```

show ospf neighbor area area-id

```
user@host >show ospf neighbor area 1.1.1.1
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.55 so-1/0/0.0          Full     10.255.245.5 128    37
Area 1.1.1.1
```

show ospf neighbor interface interface-name

```
user@host >show ospf neighbor interface so-0/0/0.0
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    37
Area 0.0.0.0
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    32
Area 2.2.2.2
```

show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled)

```
user @host > show ospf3 neighbor instance all
Instance: ina
Realm: ipv6-unicast
ID          Interface          State    Pri    Dead
100.1.1.1    fe-0/0/2.0          Full     128    37
Neighbor-address fe80::217:cb00:c87c:8c03
Instance: inb
Realm: ipv4-unicast
ID          Interface          State    Pri    Dead
100.1.2.1    fe-0/0/2.1          Full     128    33
Neighbor-address fe80::217:cb00:c97c:8c03
```

show (ospf | ospf3) overview

Syntax	show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Database protection introduced in Junos 10.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display Open Shortest Path First (OSPF) overview information.
Options	<p>none—Display standard information about all OSPF neighbors for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	show ospf overview on page 466 show ospf overview (With Database Protection) on page 467 show ospf3 overview (With Database Protection) on page 467 show ospf overview extensive on page 467
Output Fields	Table 27 on page 464 lists the output fields for the show ospf overview command. Output fields are listed in the approximate order in which they appear.

Table 27: show ospf overview Output Fields

Field name	Field Description	Level of Output
Instance	OSPF routing instance.	All levels

Table 27: show ospf overview Output Fields (*continued*)

Field name	Field Description	Level of Output
Router ID	Router ID of the routing device.	All levels
Route table index	Route table index.	All levels
Configured overload	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
Topology	Topology identifier.	All levels
Prefix export count	Number of prefixes exported into OSPF.	All levels
Full SPF runs	Number of complete Shortest Path First calculations.	All levels
SPF delay	Delay before performing consecutive Shortest Path First calculations.	All levels
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.	All levels
LSA refresh time	Refresh period for link-state advertisement (in minutes).	All levels
Database protection state	Current state of database protection.	All levels
Warning threshold	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
Non self-generated LSAs	Number of LSAs whose router ID is not equal to the local router ID: Current , Warning (threshold), and Allowed .	All levels
Ignore time	How long the database has been in the ignore state.	All levels
Reset time	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
Ignore count	Number of times the database has been in the ignore state: Current and Allowed .	All levels
Restart	Graceful restart capability: enabled or disabled .	All levels
Restart duration	Time period for complete reacquisition of OSPF neighbors.	All levels
Restart grace period	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels

Table 27: show ospf overview Output Fields (*continued*)

Field name	Field Description	Level of Output
Graceful restart helper mode	(OSPFv2) Standard graceful restart helper capability (based on RFC 3623): enabled or disabled .	All levels
Restart-signaling helper mode	(OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): enabled or disabled .	All levels
Helper mode	(OSPFv3) Graceful restart helper capability: enabled or disabled .	All levels
Trace options	OSPF-specific trace options.	extensive
Trace file	Name of the file to receive the output of the tracing operation.	extensive
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: Normal Stub , Not Stub , or Not so Stubby Stub .	All levels
Authentication Type	Type of authentication: None , Password , or MD5 . NOTE: The Authentication Type field refers to the authentication configured at the [edit protocols ospf area area-id] level. Any authentication configured for an interface in this area will not affect the value of this field.	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

Sample Output

show ospf overview

```

user@host> show ospf overview
Instance: master
  Router ID: 10.255.245.6
  Route table index: 0
  Configured overload, expires in 118 seconds
  LSA refresh time: 50 minutes
  Restart: Enabled
    Restart duration: 20 sec
    Restart grace period: 40 sec
    Helper mode: enabled
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 0
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 1

```


SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
Router ID: 10.255.112.218
Route table index: 0
LSA refresh time: 50 minutes
Traffic engineering
Restart: Enabled
  Restart duration: 180 sec
  Restart grace period: 210 sec
  Graceful restart helper mode: Enabled
  Restart-signaling helper mode: Enabled
Database protection state: Normal
Warning threshold: 70 percent
Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
Ignore time: 30, Reset time: 60
Ignore count: Current 0, Allowed 1
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 160
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 70
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed
```

show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
Router ID: 10.255.112.128
Route table index: 0
LSA refresh time: 50 minutes
Database protection state: Normal
Warning threshold: 80 percent
Non self-generated LSAs: Current 3, Warning 8, Allowed 10
Ignore time: 30, Reset time: 60
Ignore count: Current 0, Allowed 2
Area: 0.0.0.0
  Stub type: Not Stub
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 7
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed
```

show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
Router ID: 1.1.1.103
Route table index: 0
```

```
Full SPF runs: 13, SPF delay: 0.200000 sec
LSA refresh time: 50 minutes
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
```

show (ospf | ospf3) route

Syntax	<pre>show (ospf ospf3) route <brief detail extensive> <abr asbr extern inter intra> <destination> <instance (default ipv4-multicast <i>instance-name</i>)> <logical-system (default ipv4-multicast <i>logical-system-name</i>)> <network> <no-backup-coverage> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <router> <topology (default ipv4-multicast <i>topology-name</i>)> <transit></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show (ospf ospf3) route <brief detail extensive> <abr asbr extern inter intra> <destination> <instance <i>instance-name</i> <network> <no-backup-coverage> <router> <topology (default ipv4-multicast <i>topology-name</i>)> <transit></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>topology option introduced in Junos OS Release 9.0.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the Open Shortest Path First (OSPF) routing table.
Options	<p>none—Display standard information about all entries in the OSPF routing table for all routing instances and all topologies.</p> <p>destination—Display routes to the specified IP address (with optional destination prefix length).</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>abr—(Optional) Display routes to area border routers.</p> <p>asbr—(Optional) Display routes to autonomous system border routers.</p> <p>extern—(Optional) Display external routes.</p> <p>inter—(Optional) Display interarea routes.</p> <p>intra—(Optional) Display intra-area routes.</p>

instance (**default** | **ipv4-multicast** | *instance-name*)—(Optional) Display entries for the default routing instance, the IPv4 multicast routing instance, or for the specified routing instance.

logical-system (**default** | **ipv4-multicast** | *logical-system-name*)—(Optional) Perform this operation on the default logical system, the IPv4 multicast logical system, or on a particular logical system.

network—(Optional) Display routes to networks.

no-backup-coverage—(Optional) Display routes with no backup coverage.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Display entries in the routing table for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Display routes to all routers.

topology (**default** | **ipv4-multicast** | *topology-name*)—(OSPFv2 only) (Optional) Display routes for the default OSPF topology, IPv4 multicast topology, or for a particular topology.

transit—(Optional) (OSPFv3 only) Display OSPFv3 routes to pseudonodes.

Required Privilege Level view

List of Sample Output [show ospf route on page 472](#)
[show ospf route detail on page 472](#)
[show ospf3 route on page 472](#)
[show ospf3 route detail on page 473](#)
[show ospf route topology voice on page 473](#)

Output Fields [Table 28 on page 470](#) list the output fields for the **show (ospf | ospf3) route** command. Output fields are listed in the approximate order in which they appear.

Table 28: show (ospf | ospf3) route Output Fields

Field Name	Field Description	Output Level
Topology	Name of the topology.	All levels
Prefix	Destination of the route.	All levels
Path type	How the route was learned: <ul style="list-style-type: none"> Inter—Interarea route Ext1—External type 1 route Ext2—External type 2 route Intra—Intra-area route 	All levels

Table 28: show (ospf | ospf3) route Output Fields (*continued*)

Field Name	Field Description	Output Level
Route type	The type of routing device from which the route was learned: <ul style="list-style-type: none"> • AS BR—Route to AS border router. • Area BR—Route to area border router. • Area/AS BR—Route to router that is both an Area BR and AS BR. • Network—Network router. • Router—Route to a router that is neither an Area BR nor an AS BR. • Transit—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link. • Discard—Route to a summary discard. 	All levels
NH Type	Next-hop type: LSP or IP .	All levels
Metric	Route's metric value.	All levels
NH-interface	(OSPFv3 only) Interface through which the route's next hop is reachable.	All levels
NH-addr	(OSPFv3 only) IPv6 address of the next hop.	All levels
NextHop Interface	(OSPFv2 only) Interface through which the route's next hop is reachable.	All levels
Nexthop addr/label	(OSPFv2 only) If the NH Type is IP , then it is the address of the next hop. If the NH Type is LSP , then it is the name of the label-switched path.	All levels
Area	Area ID of the route.	detail
Origin	Router from which the route was learned.	detail
Type 7	Route was learned through a not-so-stubby area (NSSA) link-state advertisement (LSA).	detail
P-bit	Route was learned through NSSA LSA and the propagate bit was set.	detail
Fwd NZ	Forwarding address is nonzero. Fwd NZ is only displayed if the route is learned through an NSSA LSA.	detail
optional-capability	Optional capabilities propagated in the router LSA. This field is in the output for intra-area router routes only (when Route Type is Area BR , AS BR , Area/AS BR , or Router), not for interarea router routes or network routes. Three bits in this field are defined as follows: <ul style="list-style-type: none"> • 0x4 (V)—Routing device is at the end of a virtual active link. • 0x2 (E)—Routing device is an autonomous system boundary router. • 0x1 (B)—Routing device is an area border router. 	detail

Table 28: show (ospf | ospf3) route Output Fields (*continued*)

Field Name	Field Description	Output Level
priority	The priority assigned to the prefix: <ul style="list-style-type: none"> • high • medium • low <p>NOTE: The priority field applies only to routes of type Network.</p>	detail

Sample Output

show ospf route

```

user@host> show ospf route
Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type   Type    Type                Interface    addr/label
10.255.71.12     Intra Router  IP    1      fe-0/0/2.0   192.16.22.86
10.255.71.13/32  Intra Network IP    0      lo0.0
192.168.222.84/30 Intra Network LSP   1      fe-0/0/2.0   1sp-ab

```

show ospf route detail

```

user@host> show ospf route detail
Topology default Route Table:

Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type   Type    Type                Interface    addr/label
10.255.14.174    Inter AS BR   IP    210    t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185
10.255.14.178    Intra Router  IP    200    t3-3/1/3.0
  area 0.0.0.2, origin 10.255.14.178, optional-capability 0x0
10.210.1.0/30    Intra Network IP    10     t3-3/1/2.0
  area 0.0.0.2, origin 10.255.14.172, priority medium
100.1.1.1/32     Inter Network IP    210    t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185, priority low
112.3.1.0/24     Ext2  Network  IP    0      t1-3/0/1.0
  area 0.0.0.0, origin 10.255.14.174, priority high
200.3.3.0/30     Inter Network IP    220    t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185, priority high

```

show ospf3 route

```

user@host> show ospf3 route
Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type   Type    Type                Interface    addr/label
10.255.71.13     Intra Router  IP    1      fe-0/0/2.0   192.16.22.86
  NH-interface fe-0/0/2.0, NH-addr fe80::290:69ff:fe9b:e002
10.255.71.13;0.0.0.2
10.255.245.1     Intra Router  IP    40     fxp1.1       192.168.36.17
  area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3     Intra AS BR   IP    1      fxp2.3       192.168.36.34
  area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32  Intra Network IP    40     fxp1.1       192.168.36.17

```

```

    area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32      Intra Network   IP      0  lo0.0
    area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32      Intra Network   IP      1  fxp2.3      192.168.36.34

    area 0.0.0.0, origin 10.255.245.3, priority low
                Intra Transit   IP      1
    NH-interface fe-0/0/2.0
192::168:222:84/126 Intra Network   IP      1
    NH-interface fe-0/0/2.0
abcd::71:12/128     Intra Network   IP      0
    NH-interface lo0.0
abcd::71:13/128     Intra Network   LSP     1
    NH-interface fe-0/0/2.0, NH-addr lsp-cd

```

show ospf3 route detail

```

user@host> show ospf3 route detail
Prefix                                Path   Route   NH   Metric
                                type  type   type
10.255.14.174                        Intra Area/AS BR IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Optional-capability 0x3
10.255.14.178                        Intra Router IP    200
    NH-interface t3-3/1/3.0
    Area 0.0.0.0, Origin 10.255.14.178, Optional-capability 0x0
10.255.14.185;0.0.0.2                Intra Transit IP    200
    NH-interface t1-3/0/1.0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.185
1000:1:1::1/128                      Inter Network IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Priority low
1001:2:1::/48                        Ext1  Network IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority medium
1002:1:7::/48                        Ext2  Network IP    0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority low
1002:3:4::/48                        Ext2  Network IP    0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority high
abcd::10:255:14:172/128              Intra Network IP    0
    NH-interface lo0.0
    Area 0.0.0.0, Origin 10.255.14.172, Priority low

```

show ospf route topology voice

```

user@host show ospf route topology voice
Topology voice Route Table:
Prefix      Path   Route   NH   Metric  NextHop      Nexthop
            Type  Type   Type
10.255.8.2  Intra Router IP     1    so-0/2/0.0
10.255.8.3  Intra Router IP     2    so-0/2/0.0
10.255.8.1/32 Intra Network IP     0    lo0.0
10.255.8.2/32 Intra Network IP     1    so-0/2/0.0
10.255.8.3/32 Intra Network IP     2    so-0/2/0.0
192.168.8.0/29 Intra Network IP     2    so-0/2/0.0
192.168.8.44/30 Intra Network IP     2    so-0/2/0.0
192.168.8.46/32 Intra Network IP     1    so-0/2/0.0

```

192.168.8.48/30	Intra	Network	IP	1	so-0/2/1.0
192.168.8.52/30	Intra	Network	IP	2	so-0/2/0.0
192.168.9.44/30	Intra	Network	IP	1	so-0/2/0.0
192.168.9.45/32	Intra	Network	IP	2	so-0/2/0.0

show (ospf | ospf3) statistics

Syntax	show (ospf ospf3) statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display OSPF statistics.
Options	<p>none—Display OSPF statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display all statistics for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Display all statistics for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear (ospf ospf3) statistics on page 426
List of Sample Output	show ospf statistics on page 477 show ospf statistics logical-system all on page 477 show ospf3 statistics on page 478
Output Fields	Table 29 on page 475 lists the output fields for the show (ospf ospf3) statistics command. Output fields are listed in the approximate order in which they appear.

Table 29: show (ospf | ospf3) statistics Output Fields

Field Name	Field Description
Packet type	Type of OSPF packet.
Total Sent/Total Received	Total number of packets sent and received.
Last 5 seconds Sent/Last 5 seconds Received	Total number of packets sent and received in the last 5 seconds.

Table 29: show (ospf | ospf3) statistics Output Fields (*continued*)

Field Name	Field Description
DBDs retransmitted	Total number of database description packets retransmitted, and number retransmitted in the last 5 seconds.
LSAs flooded	Total number of link-state advertisements flooded, and number flooded in the last 5 seconds.
LSAs flooded high-prio	<p>Total number of high priority link-state advertisements flooded, and number flooded in the last 5 seconds.</p> <p>A link-state advertisement is deemed a high priority if it has changed since it was last sent.</p>
LSAs retransmitted	Total number of link-state advertisements retransmitted, and number retransmitted in the last 5 seconds.
LSAs transmitted to nbr	Total number of link-state advertisements transmitted to a neighbor, and number transmitted in the last 5 seconds.
LSAs requested	Total number of link-state advertisements requested by neighboring devices, and number requested in the last 5 seconds.
LSAs acknowledged	Total number of link-state advertisements acknowledged, and number acknowledged in the last 5 seconds.
Flood queue depth	Total number of entries in the extended queue.
Total rexmit entries	Total number of retransmission entries waiting to be sent from the OSPF routing instance.
db summaries	Total number of database description summaries waiting to be sent from the OSPF routing instance.
lsreq entries	Total number of link-state request entries waiting to be sent from the OSPF routing instance.
Receive errors	<p>Number and type of receive errors. Some sample receive errors include:</p> <ul style="list-style-type: none"> • mtu mismatches • no interface found • no virtual link found • nssa mismatches • stub area mismatches • subnet mismatches <p>If there are no receive errors, the output displays none.</p>

Sample Output

show ospf statistics

```

user@host> show ospf statistics
Packet type          Total
                   Sent      Received
Hello                31         14
  DbD                 9         10
  LSReq               2          2
LSUpdate             8         16
  LSAck              9          9
                   Last 5 seconds
                   Sent      Received
Hello                2          2
  DbD                 0          0
  LSReq               0          0
LSUpdate             0          0
  LSAck              0          0

DBDs retransmitted   :          3, last 5 seconds :          0
LSAs flooded         :         12, last 5 seconds :          0
LSAs flooded high-prio :          0, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          3, last 5 seconds :          0
LSAs requested       :          5, last 5 seconds :          0
LSAs acknowledged    :         19, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:
  862 no interface found
  115923 no virtual link found

```

show ospf statistics logical-system all

```

user@host> show ospf statistics logical-system all
logical-system: C
OSPF instance is not running
-----

logical-system: B
Packet type          Total
                   Sent      Received
Hello              313740      313653
  DbD                3          2
  LSReq              1          1
LSUpdate           2752       1825
  LSAck             1821       2747
                   Last 5 seconds
                   Sent      Received
Hello                1          0
  DbD                 0          0
  LSReq               0          0
LSUpdate            0          0
  LSAck              0          0

DBDs retransmitted   :          0, last 5 seconds :          0
LSAs flooded         :        2741, last 5 seconds :          0
LSAs flooded high-prio :         10, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          2, last 5 seconds :          0
LSAs requested       :          1, last 5 seconds :          0
LSAs acknowledged    :       1831, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:

```

```

None
-----

logical-system: A

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello                313698      313695         0         0
  DbD                  2         3         0         0
  LSReq                1         1         0         0
LSUpdate             1825      2752         0         0
LSAck                2747      1821         0         0

DBDs retransmitted   :          0, last 5 seconds :      0
LSAs flooded         :      1825, last 5 seconds :      0
LSAs flooded high-prio :      10, last 5 seconds :      0
LSAs retransmitted   :          0, last 5 seconds :      0
LSAs transmitted to nbr:        1, last 5 seconds :      0
LSAs requested       :          2, last 5 seconds :      0
LSAs acknowledged   :     2748, last 5 seconds :      0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:
None
-----

```

show ospf3 statistics

```

user@host> show ospf3 statistics

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello                0         0         0         0
  DbD                  0         0         0         0
  LSReq                0         0         0         0
LSUpdate             0         0         0         0
LSAck                0         0         0         0

DBDs retransmitted   :          0, last 5 seconds :      0
LSAs flooded         :          0, last 5 seconds :      0
LSAs flooded high-prio :          0, last 5 seconds :      0
LSAs retransmitted   :          0, last 5 seconds :      0
LSAs transmitted to nbr:          0, last 5 seconds :      0
LSAs requested       :          0, last 5 seconds :      0
LSAs acknowledged   :          0, last 5 seconds :      0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:
None

```

show as-path

Syntax	show as-path <brief detail> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show as-path <brief detail>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the as-path-prepend action) or through aggregation.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p>
Options	<p>none—Display basic information about AS paths that the local routing device is using (same as brief).</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show as-path summary on page 485
List of Sample Output	show as-path on page 480 show as-path detail on page 481
Output Fields	Table 30 on page 479 lists the output fields for the show as-path command. Output fields are listed in the approximate order in which they appear.

Table 30: show as-path Output Fields

Field Name	Field Description	Level of Output
Total AS paths	Total number of AS paths.	brief none
Bucket	Bucket number.	All levels

Table 30: show as-path Output Fields (*continued*)

Field Name	Field Description	Level of Output
Count	Number of AS path entries in this bucket.	All levels
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. • Atomic—Route is an aggregate of several route prefixes. • Aggregator—Routing device has summarized a range of prefixes. 	All levels
domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.	detail
neighbor as	AS peer address.	detail
length	Length of the AS path.	detail
segments	Length of the AS segment descriptor.	detail
references	Path reference count.	detail

Sample Output

show as-path

```

user@host> show as-path
Total AS paths: 30382
Bucket 0      Count: 36
I
14203 2914 174 31752 I
14203 2914 701 21512 I
14203 2914 1239 26632 I
14203 2914 1239 29704 I
14203 2914 4323 10248 I
14203 2914 4766 23560 I
14203 2914 6395 32776 I
14203 2914 7911 11272 I
14203 2914 12180 18440 I
14203 2914 17408 17416 I
14203 2914 701 702 24586 I
14203 2914 1239 4657 9226 I
14203 2914 1239 7132 16394 I
14203 2914 1299 8308 34826 I
14203 2914 3320 5603 28682 I
14203 2914 3491 1680 33802 I
14203 2914 3549 7908 27658 I
14203 2914 3549 20804 30730 I
14203 2914 7018 2687 9226 I
14203 2914 174 9318 9318 23564 I

```

```

14203 2914 701 3786 3786 23564 I
14203 2914 701 4761 4795 9228 I
14203 2914 1239 7132 5673 18444 I
14203 2914 3491 20485 24588 24588 I
14203 2914 5511 2200 1945 2060 I
14203 2914 7911 14325 14325 14348 I
14203 2914 701 4637 9230 9230 9230 I
14203 2914 6395 14 14 14 14 I
14203 2914 9299 6163 6163 6163 6163 9232 I
14203 2914 3356 3356 3356 3356 3356 11955 21522 I
14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
14203 2914 174 30209 30222 30222 30222 ?
14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
Bucket 1    Count: 28
14203 2914 35847 I
14203 2914 174 19465 I
14203 2914 174 35849 I
14203 2914 2828 32777 I
14203 2914 4323 14345 I
14203 2914 4323 29705 I
14203 2914 6395 32777 I

```

...

show as-path detail

```

user@host> show as-path detail
Total AS paths: 30410
Bucket 0    Count: 36
AS path: I
  domain 0, length 0, segments 0, references 54
AS path: 14203 2914 174 31752 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 701 21512 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 26632 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 29704 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4323 10248 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4766 23560 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 6395 32776 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 7911 11272 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 12180 18440 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 17408 17416 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 701 702 24586 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 1239 4657 9226 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 7
AS path: 14203 2914 1239 7132 16394 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 1299 8308 34826 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2

```

```

AS path: 14203 2914 3320 5603 28682 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3491 1680 33802 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 7908 27658 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 20804 30730 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 7018 2687 9226 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 174 9318 9318 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 3786 3786 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4761 4795 9228 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 14
AS path: 14203 2914 1239 7132 5673 18444 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 3491 20485 24588 24588 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 4
AS path: 14203 2914 5511 2200 1945 2060 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 7911 14325 14325 14348 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4637 9230 9230 9230 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 3
AS path: 14203 2914 6395 14 14 14 14 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 10
...

```


show as-path domain

Syntax	show as-path domain <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show as-path domain
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display autonomous system (AS) path domain information.
Options	none —(Optional) Display AS path domain information for all routing instances. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show as-path domain on page 484
Output Fields	Table 31 on page 483 lists the output fields for the show as-path domain command. Output fields are listed in the approximate order in which they appear

Table 31: show as-path domain Output Fields

Field Name	Field Description
Domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.
Primary	Primary AS number.
References	Path reference count.
Number Paths	Number of known AS paths.
Flags	Information about the AS path: <ul style="list-style-type: none"> • ASLoop—Path contains an AS loop. • Atomic—Path includes the ATOMIC_AGGREGATE path attribute. • Local—Path was created by local aggregation. • Master—Path was created by the master routing instance.
Local AS	AS number of the local routing device.
Loops	How many times this AS number can appear in an AS path.

Sample Output

show as-path domain

```
user@host> show as-path domain
Domain: 1          Primary: 10458
References:        3 Paths:      30383
Flags: Master
Local AS: 10458   Loops: 1
```

show as-path summary

Syntax	show as-path summary <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show as-path summary
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display autonomous system (AS) path summary information. AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.
Options	none —(Optional) Display AS path summary information for all routing instances. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show as-path on page 479
List of Sample Output	show as-path summary on page 486
Output Fields	Table 32 on page 485 lists the output fields for the show as-path summary command. Output fields are listed in the approximate order in which they appear.

Table 32: show as-path summary Output Fields

Field Name	Field Description
AS Paths	Number of AS paths.
Buckets	Number of hash buckets in use.
Max	Maximum number of AS path entries per bucket.
Min	Minimum number of AS path entries per bucket.
Avg	Average number of AS path entries per bucket.
Std deviation	Standard deviation of AS path entries per bucket.

Sample Output

show as-path summary

```
user@host> show as-path summary
AS  Paths  Buckets  Max  Min  Avg  Std deviation
30425    1024    95   12   29   6.481419
```

show bgp bmp

Syntax	show bgp bmp
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display information about the BGP Monitoring Protocol (BMP).
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show bgp bmp on page 487
Output Fields	Table 33 on page 487 lists the output fields for the show bgp bmp command. Output fields are listed in the approximate order in which they appear.

Table 33: show bgp bmp Output Fields

Field Name	Field Description
BMP station address/port	IP address and port number of the monitoring station to which BGP Monitoring Protocol (BMP) statistics are sent.
BMP session state	Status of the BMP session: UP or DOWN .
Memory consumed by BMP	Memory used by the active BMP session.
Statistics timeout	Amount of time, in seconds, between transmissions of BMP data to the monitoring station.
Memory limit	Threshold, in bytes, at which the routing device stops collecting BMP data.
Memory-connect retry timeout	Amount of time, in seconds, after which the routing device attempts to resume a BMP session that was ended after the configured memory threshold was exceeded.

Sample Output

show bgp bmp

```

user@host> show bgp bmp
  BMP station address/port: 172.24.24.157+5454
  BMP session state: DOWN
  Memory consumed by BMP: 0
  Statistics timeout: 15
  Memory limit: 10485760
  Memory connect retry timeout: 600

```


show bgp group

Syntax	<pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name> <logical-system (all logical-system-name)> <rtf></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about the configured BGP groups.
Options	<p>none—Display group information about all BGP groups.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display group information for the specified group.</p> <p>exact-instance instance-name—(Optional) Display information for the specified instance only.</p> <p>instance instance-name—(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp group instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>rtf—(Optional) Display BGP group route targeting information.</p>
Required Privilege Level	view
List of Sample Output	<p>show bgp group on page 493</p> <p>show bgp group brief on page 493</p> <p>show bgp group detail on page 494</p> <p>show bgp group rtf detail on page 495</p> <p>show bgp group summary on page 495</p>

Output Fields Table 34 on page 490 describes the output fields for the **show bgp group** command. Output fields are listed in the approximate order in which they appear.

Table 34: show bgp group Output Fields

Field Name	Field Description	Level of Output
Group Type or Group	Type of BGP group: Internal or External .	All levels
group-index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.	rtf detail
AS	AS number of the peer. For internal BGP (IBGP), this number is the same as Local AS .	brief detail none
Local AS	AS number of the local routing device.	brief detail none
Name	Name of a specific BGP group.	brief detail none
Index	Unique index number of a BGP group.	brief detail none
Flags	Flags associated with the BGP group. This field is used by Juniper Networks customer support.	brief detail none
Remove-private options	Options associated with the remove-private statement.	brief detail none
Holdtime	Maximum number of seconds allowed to elapse between successive keepalive or update messages that BGP receives from a peer in the BGP group, after which the connection to the peer is closed and routing devices through that peer become unavailable.	brief detail none
Export	Export policies configured for the BGP group with the export statement.	brief detail none
MED tracks IGP metric update delay	Time, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire	All levels
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.	brief detail none
Total peers	Total number of peers in the group.	brief detail none
Established	Number of peers in the group that are in the established state.	All levels

Table 34: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active/Received/Accepted/Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established in the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. 	summary
ip-addresses	List of peers who are members of the group. The address is followed by the peer's port number.	All levels
Route Queue Timer	Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed.	detail
Route Queue	Number of prefixes that are queued up for sending to the peers in the group.	detail
inet.number	<p>Number of active, received, accepted, and damped routes in the routing table. For example, inet.0: 7/10/9/0 indicates the following:</p> <ul style="list-style-type: none"> 7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the inet.0 routing table. 	none

Table 34: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Table inet.number	Information about the routing table. <ul style="list-style-type: none"> • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. • Advertised prefixes—Number of prefixes advertised to a peer. • Received external prefixes—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table. • Active external prefixes—Number of prefixes received from the EBGP peers that are active in the routing table. • Externals suppressed—Number of routes received from EBGP peers currently inactive because of damping or other reasons. • Received internal prefixes—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table. • Active internal prefixes—Number of prefixes received from the IBGP peers that are active in the routing table. • Internals suppressed—Number of routes received from IBGP peers currently inactive because of damping or other reasons. • RIB State—Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete. 	detail
Groups	Total number of groups.	All levels
Peers	Total number of peers.	All levels
External	Total number of external peers.	All levels
Internal	Total number of internal peers.	All levels
Down peers	Total number of unavailable peers.	All levels
Flaps	Total number of flaps that occurred.	All levels
Table	Name of a routing table.	brief , none
Tot Paths	Total number of routes.	brief , none
Act Paths	Number of active routes.	brief , none
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	brief , none

Table 34: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
History	Number of withdrawn routes stored locally to keep track of damping history.	brief, none
Damp State	Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs.	brief, none
Pending	Routes being processed by the BGP import policy.	brief, none
Group	Group the peer belongs to in the BGP configuration.	detail
Receive mask	Mask of the received target included in the advertised route.	detail
Entries	Number of route entries received.	detail
Target	Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer.	detail
Mask	Mask which specifies that the peer receive routes with the given route target.	detail

Sample Output

show bgp group

```

user@host> show bgp group
Groups: 2  Peers: 2   External: 0   Internal: 2   Down peers: 1   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed   History Damp State   Pending

inet.0
          0         0         0         0         0         0

bgp.13vpn.0
          0         0         0         0         0         0

bgp.rtarget.0
          2         0         0         0         0         0

```

show bgp group brief

```

user@host> show bgp group brief
Groups: 2  Peers: 2   External: 0   Internal: 2   Down peers: 1   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed   History Damp State   Pending

inet.0
          0         0         0         0         0         0

bgp.13vpn.0
          0         0         0         0         0         0

```

```

bgp.rtarget.0
                2          0          0          0          0          0

```

show bgp group detail

```

user@host> show bgp group detail
Group Type: Internal  AS: 1                      Local AS: 1
Name: ibgp           Index: 0                  Flags: <Export Eval>
Holdtime: 0
Total peers: 3       Established: 0
22.0.0.2
22.0.0.8
22.0.0.5

Groups: 1  Peers: 3  External: 0  Internal: 3  Down peers: 3  Flaps: 3
Table bgp.l3vpn.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table bgp.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.inet.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0

```

```

Externals suppressed:      0
Received internal prefixes: 0
Active internal prefixes:  0
Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete

```

show bgp group rtf detail

```

user@host> show bgp group rtf detail
Group: internal (group-index: 0)
  Receive mask: 00000002
  Table: bgp.rtarget.0
    Target      Mask      Entries: 2
    100:100/64  00000002
    200:201/64  (Group)
Group: internal (group-index: 1)
  Table: bgp.rtarget.0
    Target      Mask      Entries: 1
    200:201/64  (Group)

```

show bgp group summary

```

user@host> show bgp group summary
Group      Type      Peers    Established    Active/Received/Accepted/Damped
ibgp       Internal  3        0

```

Groups: 1 Peers: 3 External: 0 Internal: 3 Down peers: 3 Flaps: 3

```

bgp.l3vpn.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
bgp.mdt.0   : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.inet.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.mdt.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0

```

show bgp neighbor

Syntax	<pre>show bgp neighbor <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor-address> <orf (detail <i>neighbor-address</i>)</pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp neighbor <instance <i>instance-name</i>> <exact-instance <i>instance-name</i>> <neighbor-address> <orf (<i>neighbor-address</i> detail)</pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>orf option introduced in Junos OS Release 9.2.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about BGP peers.
Options	<p>none—Display information about all BGP peers.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp neighbor instance cust1 command).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p>orf (detail <i>neighbor-address</i>)—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the detail option to display detailed output.</p>
Additional Information	For information about the local-address , nlri , hold-time , and preference statements, see the Junos OS Routing Protocols Configuration Guide.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear bgp neighbor on page 429

List of Sample Output [show bgp neighbor on page 503](#)
[show bgp neighbor \(CLNS\) on page 504](#)
[show bgp neighbor \(Layer 2 VPN\) on page 504](#)
[show bgp neighbor \(Layer 3 VPN\) on page 506](#)
[show bgp neighbor neighbor-address on page 507](#)
[show bgp neighbor neighbor-address on page 508](#)
[show bgp neighbor orf neighbor-address detail on page 509](#)

Output Fields Table 35 on page 497 describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 35: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor port number.
AS	AS number of the peer.
Local	Address of the local routing device. The address is followed by the peer port number.
Type	Type of peer: Internal or External .
State	<p>Current state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Flags	<p>Internal BGP flags:</p> <ul style="list-style-type: none"> • Aggregate Label—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • CleanUp—The peer session is being shut down. • Delete—This peer has been deleted. • Idled—This peer has been permanently idled. • ImportEval—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • Initializing—The peer session is initializing. • SendRtn—Messages are being sent to the peer. • Sync—This peer is synchronized with the rest of the peer group. • TryConnect—Another attempt is being made to connect to the peer. • Unconfigured—This peer is not configured. • WriteFailed—An attempt to write to this peer failed.

Table 35: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Last state	<p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Last event	<p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Closed—The BGP session closed. • ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect. • HoldTime—The session ended because the hold timer expired. • KeepAlive—The local routing device sent a BGP keepalive message to the peer. • Open—The local routing device sent a BGP open message to the peer. • OpenFail—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • RecvKeepAlive—The local routing device received a BGP keepalive message from the peer. • RecvNotify—The local routing device received a BGP notification message from the peer. • RecvOpen—The local routing device received a BGP open message from the peer. • RecvUpdate—The local routing device received a BGP update message from the peer. • Start—The peering session started. • Stop—The peering session stopped. • TransportError—A TCP error occurred.
Last error	<p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Cease—An error occurred, such as a version mismatch, that caused the session to close. • Finite State Machine Error—In setting up the session, BGP received a message that it did not understand. • Hold Time Expired—The session's hold time expired. • Message Header Error—The header of a BGP message was malformed. • Open Message Error—A BGP open message contained an error. • None—No errors occurred in the BGP session. • Update Message Error—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.

Table 35: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Options	Configured BGP options: <ul style="list-style-type: none"> • AddressFamily—Configured address family: inet or inet-vpn. • AuthKeyChain—Authentication key change is enabled. • DropPathAttributes—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • GracefulRestart—Graceful restart is configured. • HoldTime—Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent. • IgnorePathAttributes—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • Local Address—Address configured with the local-address statement. • Multihop—Allow BGP connections to external peers that are not on a directly connected network. • NLRI—Configured MBGP state for the BGP group: multicast, unicast, or both if you have configured nlri any. • Peer AS—Configured peer autonomous system (AS). • Preference—Preference value configured with the preference statement. • Refresh—Configured to refresh automatically when the policy changes. • Rib-group—Configured routing table group.
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Authentication key change	(appears only if the authentication-keychain statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(appears only if the authentication-algorithm statement has been configured) Type of authentication algorithm enabled: hmac or md5 .
Address families configured	Names of configured address families for the VPN.
Local Address	Address of the local routing device.
Remove-private options	Options associated with the remove-private statement.
Holdtime	Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: <ul style="list-style-type: none"> • TrafficStatistics—Collection of statistics for labeled-unicast traffic is enabled.

Table 35: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Traffic statistics	Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> • Options—Options configured for collecting statistics about labeled-unicast traffic. • File—Name and location of statistics log files. • size—Size of all the log files, in bytes. • files—Number of log files.
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the preference statement.
Outbound Timer	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the out-delay parameter is configured to a non-zero value.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.
Group index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.
Peer index	Index that is unique within the BGP group to which the peer belongs.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time that the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGP peering is established.
NLRI for restart configured on peer	Names of address families configured for restart.
NLRI advertised by peer	Address families supported by the peer: unicast or multicast .
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .
Restart time configured on peer	Configured time allowed for restart on the neighbor.

Table 35: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Peer does not support Restarter functionality	Graceful restart restarter-mode is disabled on the peer.
Peer does not support Receiver functionality	Graceful restart helper-mode is disabled on the peer.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the end-of-rib marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Peer supports 4 byte AS extension (peer-as 1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.
NLRIs for which peer can receive multiple paths	Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route. Possible value is inet-unicast .
NLRIs for which peer can send multiple paths: inet-unicast	Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route. Possible value is inet-unicast .

Table 35: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Table inet. <i>number</i>	<p>Information about the routing table:</p> <ul style="list-style-type: none"> • RIB State—BGP is in the graceful restart process for this routing table: restart is complete or restart in progress. • Bit—Number that represents the entry in the routing table for this peer. • Send state—State of the BGP group: in sync, not in sync, or not advertising. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Accepted prefixes—Total number of prefixes from the peer that have been accepted by a routing policy. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Input dropped path attributes	<p>Information about dropped path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Input ignored path attributes	<p>Information about ignored path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Output queue	Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.
Trace options	Configured tracing of BGP protocol packets and operations.
Trace file	Name of the file to receive the output of the tracing operation.
Filter Updates rcv	<p>(orf option only) Number of outbound-route filters received for each configured address family.</p> <p>NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.</p>

Table 35: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Immediate	(orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Filter	(orf option only) Type of prefix filter received: prefix-based or extended-community .
Received filter entries	(orf option only) List of received filters displayed.
seq	(orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(orf option only) Address for the prefix entry that matches the filter.
minlength	(orf option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(orf option only) Maximum prefix length, in bits, required to match this prefix.
match	(orf option only) For this prefix match, whether to permit or deny route updates.

Sample Output

show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250   Local ID: 10.255.7.248   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 10)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:           1
    Received prefixes:         1

```

```

Accepted prefixes:          1
Suppressed due to damping:  0
Advertised prefixes:       1
Last traffic (seconds): Received 9    Sent 5    Checked 5
Input messages:  Total 36    Updates 2    Refreshes 0    Octets 718
Output messages: Total 37    Updates 1    Refreshes 0    Octets 796
Output Queue[0]: 0

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
Type: Internal    State: Established (route reflector client)Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
Address families configured: inet-unicast inet-vpn-unicast route-target
Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.255.162.214    Local ID: 10.255.167.205    Active Holdtime: 90
Keepalive Interval: 30      Group index: 0    Peer index: 1

```

show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
Type: External    State: Established    Flags: <ImportEval Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
Rib-group Refresh>
Address families configured: iso-vpn-unicast
Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.245.245.1    Local ID: 10.245.245.3    Active Holdtime: 90
Keepalive Interval: 30    Peer index: 0
NLRI advertised by peer: iso-vpn-unicast
NLRI for this session: iso-vpn-unicast
Peer supports Refresh capability (2)
Table bgp.isovpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table aaa.iso.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
Last traffic (seconds): Received 6    Sent 5    Checked 5
Input messages:  Total 1736    Updates 4    Refreshes 0    Octets 33385
Output messages: Total 1738    Updates 3    Refreshes 0    Octets 33305
Output Queue[0]: 0
Output Queue[1]: 0

```

show bgp neighbor (Layer 2 VPN)

```

user@host> show bgp neighbor
Peer: 10.69.103.2    AS 65100 Local: 10.69.103.1    AS 65103
Type: External    State: Active    Flags: <ImportEval>

```

```

Last State: Idle           Last Event: Start
Last Error: None
Export: [ BGP-INET-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-unicast
Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.69.104.2          AS 65100 Local: 10.69.104.1          AS 65104
Type: External            State: Active            Flags: <ImportEval>
Last State: Idle          Last Event: Start
Last Error: None
Export: [ BGP-L-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-labeled-unicast
Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.255.14.182+179 AS 69   Local: 10.255.14.176+2131 AS 69
Type: Internal            State: Established      Flags: <ImportEval>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast l2vpn
Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.255.14.182      Local ID: 10.255.14.176      Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2

```

```

    Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

```

show bgp neighbor (Layer 3 VPN)

```

user@host> show bgp neighbor
Peer: 4.4.4.4+179    AS 10045 Local: 5.5.5.5+1214    AS 10045
Type: Internal    State: Established    Flags: <ImportEval>

```



```

Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Export: [ match-all ] Import: [ match-all ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
        Rib-group Refresh>
Address families configured: inet-vpn-unicast
Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
Flags for NLRI inet-labeled-unicast: TrafficStatistics
Traffic Statistics: Options: all File: /var/log/bstat.log
                        size 131072 files 10

Traffic Statistics Interval: 60
Number of flaps: 0
Peer ID: 192.168.1.110    Local ID: 192.168.1.111    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast
NLRI advertised by peer: inet-vpn-unicast
NLRI for this session: inet-vpn-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast
NLRI peer can save forwarding state: inet-vpn-unicast
NLRI that peer saved forwarding for: inet-vpn-unicast
NLRI that restart is negotiated for: inet-vpn-unicast
NLRI of received end-of-rib markers: inet-vpn-unicast
NLRI of all end-of-rib markers sent: inet-vpn-unicast
Table bgp.13vpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table vpn-green.inet.0 Bit: 20001
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Last traffic (seconds): Received 15   Sent 20   Checked 20
Input messages: Total 40   Updates 2   Refreshes 0   Octets 856
Output messages: Total 44   Updates 2   Refreshes 0   Octets 1066
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr.log size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
Type: Internal State: Established (route reflector client)Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
Refresh>
Address families configured: inet-vpn-unicast inet-labeled-unicast
Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
Flags for NLRI inet-vpn-unicast: AggregateLabel

```

```

Flags for NLRI inet-labeled-unicast: AggregateLabel
Number of flaps: 0
Peer ID: 10.255.245.12    Local ID: 10.255.245.13    Active Holdtime: 90
Keepalive Interval: 30
BFD: disabled
NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
NLRI for this session: inet-vpn-unicast inet-labeled-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 300
Stale routes from peer are kept for: 60
Restart time requested by this peer: 300
NLRI that peer supports restart for: inet-unicast inet6-unicast
NLRI that restart is negotiated for: inet-unicast inet6-unicast
NLRI of received end-of-rib markers: inet-unicast inet6-unicast
NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 4
  Received prefixes: 6
  Suppressed due to damping: 0
Table inet6.0 Bit: 20000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 2
  Suppressed due to damping: 0
Last traffic (seconds): Received 3    Sent 3    Checked 3
Input messages: Total 9    Updates 6    Refreshes 0    Octets 403
Output messages: Total 7    Updates 3    Refreshes 0    Octets 365
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpgr size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
Type: External    State: Established    Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: Cease
Export: [ export-policy ] Import: [ import-policy ]
Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
Address families configured: inet-unicast inet-multicast
Holdtime: 60000 Preference: 170
Number of flaps: 4
Last flap event: RecvUpdate
Error: 'Cease' Sent: 5 Recv: 0
Peer ID: 10.255.245.6    Local ID: 10.255.245.5    Active Holdtime: 60000
Keepalive Interval: 20000    Peer index: 0
BFD: disabled, down
Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:      8
  Received prefixes:    10

```

```

Accepted prefixes:          10
Suppressed due to damping:  0
Advertised prefixes:        3
Table inet.2 Bit: 20000
RIB State: BGP restart is complete
Send state: in sync
Active prefixes:            0
Received prefixes:          0
Accepted prefixes:          0
Suppressed due to damping:  0
Advertised prefixes:        0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

show bgp neighbor orf neighbor-address detail

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
  Filter updates rcv:          1 Immediate:          1
  Filter: prefix-based receive
  Received filter entries:
    seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast
  Filter updates rcv:          0 Immediate:          1
  Filter: prefix-based receive
  Received filter entries:
    *.*

```

show bgp summary

Syntax	<pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display BGP summary information.
Options	<p>none—Display BGP summary information for all routing instances.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp summary instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	<p>show bgp summary (When a Peer Is Not Established) on page 513</p> <p>show bgp summary (When a Peer Is Established) on page 513</p> <p>show bgp summary (CLNS) on page 513</p> <p>show bgp summary (Layer 2 VPN) on page 513</p> <p>show bgp summary (Layer 3 VPN) on page 514</p>
Output Fields	<p>Table 36 on page 510 describes the output fields for the show bgp summary command. Output fields are listed in the approximate order in which they appear.</p>

Table 36: show bgp summary Output Fields

Field Name	Field Description
Groups	Number of BGP groups.
Peers	Number of BGP peers.

Table 36: show bgp summary Output Fields (*continued*)

Field Name	Field Description
Down peers	Number of down BGP peers.
Table	Name of routing table.
Tot Paths	Total number of paths.
Act Paths	Number of active routes.
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
History	Number of withdrawn routes stored locally to keep track of damping history.
Damp State	Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs.
Pending	Routes in process by BGP import policy.
Peer	Address of each BGP peer. Each peer has one line of output.
AS	Peer's AS number.
InPkt	Number of packets received from the peer.
OutPkt	Number of packets sent to the peer.
OutQ	Number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.
Flaps	Number of times the BGP session has gone down and then come back up.
Last Up/Down	Last time since the neighbor transitioned to or from the established state.

Table 36: show bgp summary Output Fields (*continued*)

Field Name	Field Description
State #Active /Received/Accepted /Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established on the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. In general, the Idle state is the first stage of a connection. BGP is waiting for a Start event. A session can be idle for other reasons as well. The reason that a session is idle is sometimes displayed. For example: Idle (Removal in progress) or Idle (LicenseFailure). If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. If a BGP session is established in a routing instance, the field indicates the established (Establ) state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, Establ VPN-AB.inet.0: 2/4/0 indicates the following: <ul style="list-style-type: none"> The BGP session is established. Routes are received in the VPN-AB.inet.0 routing table. The local routing device has two active routes, four received routes, and no damped routes from a BGP peer. <p>When a BGP session is established, the peers are exchanging update messages.</p>

Sample Output

show bgp summary (When a Peer Is Not Established)

```

user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 6 4 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.3 65002 86 90 0 2 42:54 0/0/0

0/0/0
10.0.0.4 65002 90 91 0 1 42:54 0/2/0

0/0/0
10.0.0.6 65002 87 90 0 3 3 Active
10.1.12.1 65001 89 89 0 1 42:54 4/4/0

0/0/0

```

show bgp summary (When a Peer Is Established)

```

user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 6 4 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2 65002 88675 88652 0 2 42:38 2/4/0

0/0/0
10.0.0.3 65002 54528 54532 0 1 2w4d22h 0/0/0

0/0/0
10.0.0.4 65002 51597 51584 0 0 2w3d22h 2/2/0

0/0/0

```

show bgp summary (CLNS)

```

user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.245.245.1 200 1735 1737 0 0 14:26:12 Establ
  bgp.isovpn.0: 3/3/0
  aaaa.iso.0: 3/3/0

```

show bgp summary (Layer 2 VPN)

```

user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l2vpn.0 1 1 0 0 0 0
inet.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last
Up/Dwn State|#Active/Received/Damped...
10.255.245.35 65299 72 74 0 1 19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0

```

```

10.255.245.36 65299      2164      2423      0        4        19:50 Establ
  bgp.12vpn.0: 0/0/0
  frame-vpn.12vpn.0: 0/0/0
10.255.245.37 65299      36         37         0        4        17:07 Establ
  inet.0: 0/0/0
10.255.245.39 65299      138        168         0        6        53:48 Establ
  bgp.12vpn.0: 0/0/0
  frame-vpn.12vpn.0: 0/0/0
10.255.245.69 65299      134        140         0        6        53:42 Establ
  inet.0: 0/0/0

```

show bgp summary (Layer 3 VPN)

```

user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State Pending
bgp.13vpn.0      2          2          0          0      0      0      0
Peer          AS      InPkt      OutPkt      OutQ      Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5      2          21         22          0          0      6:26 Establ
  VPN-AB.inet.0: 1/1/0
10.255.71.15   1          19         21          0          0      6:17 Establ
  bgp.13vpn.0: 2/2/0
  VPN-A.inet.0: 1/1/0
  VPN-AB.inet.0: 2/2/0
  VPN-B.inet.0: 1/1/0

```


show ipv6 neighbors

Syntax	show ipv6 neighbors
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Display information about the IPv6 neighbor cache.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ipv6 neighbors on page 433
List of Sample Output	show ipv6 neighbors on page 515
Output Fields	Table 37 on page 515 describes the output fields for the show ipv6 neighbors command. Output fields are listed in the approximate order in which they appear.

Table 37: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: up , down , incomplete , reachable , stale , or unreachable .
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: yes or no .
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: yes or no .
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```

user@host> show ipv6 neighbors
IPv6 Address          Linklayer Address  State      Exp Rtr Secure
Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale      546 yes no
fe-1/2/0.1

```

fe80::2a0:a514:0:24c fe-1/2/0.1	00:05:85:8f:c8:bd	stale	258	yes	no
fe80::2a0:a514:0:64c fe-1/2/1.5	00:05:85:8f:c8:bd	stale	111	yes	no
fe80::2a0:a514:0:a4c fe-1/2/2.9	00:05:85:8f:c8:bd	stale	327	yes	no

show isis adjacency

Syntax	<pre>show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>	
Syntax (EX Series Switches and QFX Series)	<pre>show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i>></pre>	
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>	
Description	Display information about IS-IS neighbors.	
Options	<p>none—Display standard information about IS-IS neighbors for all routing instances.</p> <p><i>system id</i>—(Optional) Display information about IS-IS neighbors for the specified intermediate system.</p> <p>brief detail extensive—(Optional) Display standard information about IS-IS neighbors with the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display information about IS-IS neighbors for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display information about IS-IS neighbors for all logical systems or for a particular logical system.</p>	
Required Privilege Level	view	
Related Documentation	<ul style="list-style-type: none"> • clear isis adjacency on page 434 	
List of Sample Output	<p>show isis adjacency on page 519</p> <p>show isis adjacency brief on page 519</p> <p>show isis adjacency detail on page 520</p> <p>show isis adjacency extensive on page 520</p>	
Output Fields	<p>Table 38 on page 517 describes the output fields for the show isis adjacency command. Output fields are listed in the approximate order in which they appear.</p>	

Table 38: show isis adjacency Output Fields

Field Name	Field Description	Level of Output
Interface	Interface through which the neighbor is reachable.	All levels

Table 38: show isis adjacency Output Fields (*continued*)

Field Name	Field Description	Level of Output
System	System identifier (sysid), displayed as a name, if possible.	brief
L or Level	Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address.	All levels
State	State of the adjacency: Up , Down , New , One-way , Initializing , or Rejected .	All levels
Hold (secs)	Remaining hold time of the adjacency.	brief
SNPA	Subnetwork point of attachment (MAC address of the next hop).	brief
Expires in	How long until the adjacency expires, in seconds.	detail
Priority	Priority to become the designated intermediate system.	detail extensive
Up/Down transitions	Count of adjacency status changes from Up to Down or from Down to Up .	detail
Last transition	Time of the last Up/Down transition.	detail
Circuit type	Bit mask of levels on this interface: 1=Level 1 router; 2=Level 2 router; 3=both Level 1 and Level 2 router.	detail
Speaks	Protocols supported by this neighbor.	detail extensive
MAC address	MAC address of the interface.	detail extensive
Topologies	Supported topologies.	detail extensive
Restart capable	Whether a neighbor is capable of graceful restart: Yes or No .	detail extensive
Adjacency advertisement: Advertise	This router has signaled to advertise this interface to its neighbors in their link-state PDUs.	detail extensive
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise the interface in the router's outbound link-state PDUs.	detail extensive
IP addresses	IP address of this neighbor.	detail extensive

Table 38: show isis adjacency Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transition log	<p>List of recent transitions, including:</p> <ul style="list-style-type: none"> • When—Time at which an IS-IS adjacency transition occurred. • State—Current state of the IS-IS adjacency (up, down, or rejected). <ul style="list-style-type: none"> • Up—Adjacency is up and operational. • Down—Adjacency is down and not available. • Rejected—Adjacency has been rejected. • Event—Type of transition that occurred. <ul style="list-style-type: none"> • Seenself—Possible routing loop has been detected. • Interface down—IS-IS interface has gone down and is no longer available. • Error—Adjacency error. • Down reason—Reason that an IS-IS adjacency is down: <ul style="list-style-type: none"> • 3-Way Handshake Failed—Connection establishment failed. • Address Mismatch—Address mismatch caused link failure. • Aged Out—Link expired. • ISO Area Mismatch—IS-IS area mismatch caused link failure. • Bad Hello—Unacceptable hello message caused link failure. • BFD Session Down—Bidirectional failure detection caused link failure. • Interface Disabled—IS-IS interface is disabled. • Interface Down—IS-IS interface is unavailable. • Interface Level Disabled—IS-IS level is disabled. • Level Changed—IS-IS level has changed on the adjacency. • Level Mismatch—Levels on adjacency are not compatible. • MPLS LSP Down—Label-switched path (LSP) is unavailable. • MT Topology Changed—IS-IS topology has changed. • MT Topology Mismatch—IS-IS topology is mismatched. • Remote System ID Changed—Adjacency peer system ID changed. • Protocol Shutdown—IS-IS protocol is disabled. • CLI Command—Adjacency brought down by user. • Unknown—Unknown. 	extensive

Sample Output

show isis adjacency

```

user@host> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
at-2/3/0.0         ranier      3 Up          23

```

show isis adjacency brief

The output for the **show isis adjacency brief** command is identical to that for the **show isis adjacency** command. For sample output, see [show isis adjacency on page 519](#).

show isis adjacency detail

```

user@host> show isis adjacency detail
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 21 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:09 ago
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 11.1.1.2
  
```

show isis adjacency extensive

```

user@host> show isis adjacency extensive
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 22 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:16 ago
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 11.1.1.2
  Transition log:
  
```

When	State	Event	Down reason
Wed Nov 8 21:24:25	Up	SeenseIf	

show isis authentication

Syntax	show isis authentication <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis authentication <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 7.5. Command introduced in Junos OS Release 9.0 for EX Series switches. Support for hitless authentication key rollover introduced in Junos OS Release 11.2. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about IS-IS authentication.
Options	<p>none—Display information about IS-IS authentication.</p> <p>instance <i>instance-name</i>—(Optional) Display IS-IS authentication for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show isis authentication on page 522 show isis authentication (With Hitless Authentication Key Rollover Configured) on page 522
Output Fields	Table 39 on page 521 describes the output fields for the show isis authentication command. Output fields are listed in the approximate order in which they appear.

Table 39: show isis authentication Output Fields

Field Name	Field Description
Interface	Interface name.
Level	IS-IS level.
IIH Auth	IS-IS Hello (IIH) packet authentication type. Displays the name of the active keychain if hitless authentication key rollover is configured.
CSN Auth	Complete sequence number authentication type.
PSN Auth	Partial sequence number authentication type.

Table 39: show isis authentication Output Fields (*continued*)

Field Name	Field Description
L1 LSP Authentication	Layer 1 link-state PDU authentication type.
L2 LSP Authentication	Layer 2 link-state PDU authentication type.

Sample Output

show isis authentication

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
at-2/3/0.0         1      Simple    Simple    Simple
                   2      MD5       MD5       MD5

L1 LSP Authentication: Simple
L2 LSP Authentication: MD5

```

show isis authentication (With Hitless Authentication Key Rollover Configured)

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
so-0/1/3.0         2      hakrhello MD5       MD5

L2 LSP Authentication: MD5

```


show isis backup coverage

Syntax	<pre>show isis backup coverage <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis backup coverage <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display information about the level of backup coverage available.
Options	<p>none—Display information about the level of backup coverage available for all the nodes and prefixes in the network.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the level of backup coverage for a specific IS-IS routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Link and Node Protection for IS-IS Routes show isis backup label-switched-path on page 525
List of Sample Output	show isis backup coverage on page 524
Output Fields	<p>Table 40 on page 523 lists the output fields for the show isis backup coverage command. Output fields are listed in the approximate order in which they appear.</p>

Table 40: show isis backup coverage Output Fields

Field Name	Field Description
Topology	Type of topology or address family: IPv4 Unicast or IPv6 Unicast .
Level	IS-IS level: <ul style="list-style-type: none"> 1—Level 1 2—Level 2
Node	By topology, the percentage of all routes configured on the node that are protected through backup coverage.

Table 40: show isis backup coverage Output Fields (*continued*)

Field Name	Field Description
IPv4	Percentage of IPv4 unicast routes that are protected through backup coverage.
IPv6	Percentage of IPv6 unicast routes that are protected through backup coverage.
CLNS	Percentage of Connectionless Network Service (CLNS) routes that are protected through backup coverage.

Sample Output

show isis backup coverage

```
user@host> show isis backup coverage
Backup Coverage:
  Topology   Level  Node   IPv4   IPv6   CLNS
  IPv4 Unicast    2  28.57%  22.22%  0.00%  0.00%
  IPV6 Unicast    2   0.00%   0.00%  0.00%  0.00%
```

show isis backup label-switched-path

Syntax	show isis backup label-switched-path <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis backup label-switched-path
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes.
Options	none —Display information about MPLS LSPs designated as backup routes for IS-IS routes. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Link and Node Protection for IS-IS Routes show isis backup coverage on page 523
List of Sample Output	show isis backup label-switched-path on page 526
Output Fields	Table 41 on page 525 lists the output fields for the show isis backup label-switched-path command. Output fields are listed in the approximate order in which they appear.

Table 41: show isis backup label-switched-path Output Fields

Field Name	Field Description
Backup MPLS LSPs	List of MPLS LSPs designated as backup paths for IS-IS routes.
Egress	IP address of the egress routing device for the LSP.
Status	State of the LSP: <ul style="list-style-type: none"> Up—The router can detect RSVP hello messages from the neighbor. Down—The router has received one of the following indications: <ul style="list-style-type: none"> Communication failure from the neighbor. Communication from IGP that the neighbor is unavailable. Change in the sequence numbers in the RSVP hello messages sent by the neighbor. Deleted—LSP is no longer available as a backup path.

Table 41: show isis backup label-switched-path Output Fields (*continued*)

Field Name	Field Description
Last change	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <i>hh:mm:ss</i> .
TE-metric	Configured traffic engineering metric.
Metric	Configured metric.

Sample Output

show isis backup label-switched-path

```
user@host> show isis backup label-switched-path
Backup MPLS LSPs:
f-to-g, Egress: 192.168.1.4, Status: up, Last change: 06:12:03
TE-metric: 9, Metric: 0
```

show isis backup spf results

Syntax	<pre>show isis backup spf results <instance <i>instance-name</i>> <level (1 2)> <logical-system (all <i>logical-system-name</i>)> <no-coverage> <topology (ipv4-unicast ipv6-multicast ipv6-unicast unicast)></pre>
Syntax (EX Series Switches)	<pre>show isis backup spf results <instance <i>instance-name</i>> <level (1 2)> <no-coverage> <topology (ipv4-unicast unicast)></pre>
Release Information	Command introduced in Junos OS Release 9.5.
Description	Display information about IS-IS shortest-path-first (SPF) calculations for backup paths.
Options	<p>none—Display information about IS-IS SPF calculations for all backup paths for all destination nodes.</p> <p>instance <i>instance-name</i>—(Optional) Display SPF calculations for backup paths for the specified routing instance.</p> <p>level (1 2)—(Optional) Display SPF calculations for the backup paths for the specified IS-IS level.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display SPF calculations for the backup paths for all logical systems or on a particular logical system.</p> <p>no-coverage—(Optional) Display SPF calculations only for destinations that do not have backup coverage.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display SPF calculations for backup paths for the specified topology only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Link and Node Protection for IS-IS Routes show isis backup coverage on page 523
List of Sample Output	<p>show isis backup spf results on page 528</p> <p>show isis backup spf results no-coverage on page 529</p>
Output Fields	<p>Table 42 on page 528 lists the output fields for the show isis backup spf results command. Output fields are listed in the approximate order in which they appear.</p>

Table 42: show isis backup spf results Output Fields

Field Name	Field Description
<i>node-name</i>	Name of the destination node.
Address	Address of the destination node.
Primary next-hop	Interface and name of the node of the primary next hop to reach the destination.
Root	Name of the next-hop neighbor.
Metric	Metric to the node.
Eligible	Indicates that the next-hop neighbor has been designated as a backup path to the destination node.
Backup next-hop	Name of the interface of the backup next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
LSP	Name of the MPLS label-switched path (LSP) designated as a backup path.
Not eligible	Indicates that the next-hop neighbor cannot function as a backup path to the destination.
Reason	Describes why the next-hop neighbor is designated as Not eligible as a backup path.

Sample Output

show isis backup spf results

```

user@host> show isis backup spf results

IS-IS level 1 SPF results:
  0 nodes

IS-IS level 2 SPF results:
banff.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
  Primary next-hop: so-6/0/0.0, IPV6, olympic
  Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
  Root: camaro, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop multipath
  Root: olympic, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop multipath
  Root: glacier, Root Metric: 10, Metric: 25
  Not eligible, Reason: Primary next-hop multipath
crater.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: so-6/0/0.0, IPV6, olympic

```

```

Root: olympic, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop link fate sharing
Root: glacier, Root Metric: 10, Metric: 15
  Eligible, Backup next-hop: as0.0, IPV4, glacier
  Eligible, Backup next-hop: as0.0, IPV6, glacier
Root: camaro, Root Metric: 10, Metric: 20
  Not eligible, Reason: Interface is already covered
olympic.00
Primary next-hop: so-6/0/0.0, IPV4, olympic
Primary next-hop: so-6/0/0.0, IPV6, olympic
Root: olympic, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: camaro, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: glacier, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
camaro.00
Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
Root: camaro, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: glacier, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: olympic, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
glacier.00
Primary next-hop: as0.0, IPV4, glacier
Primary next-hop: as0.0, IPV6, glacier
Root: glacier, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: camaro, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: olympic, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
5 nodes

```

show isis backup spf results no-coverage

```
user@host> show isis backup spf results no-coverage
```

```
IS-IS level 1 SPF results:
0 nodes
```

```
IS-IS level 2 SPF results:
olympic.00
Primary next-hop: so-6/0/0.0, IPV4, olympic
Primary next-hop: so-6/0/0.0, IPV6, olympic
Root: olympic, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing

```

```

Root: camaro, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: glacier, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
camaro.00
  Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
  Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
  Root: camaro, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
  Root: glacier, Root Metric: 10, Metric: 20
    track-item: camaro.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
  Root: olympic, Root Metric: 10, Metric: 20
    track-item: camaro.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
glacier.00
  Primary next-hop: as0.0, IPV4, glacier
  Primary next-hop: as0.0, IPV6, glacier
  Root: glacier, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
  Root: camaro, Root Metric: 10, Metric: 20
    track-item: glacier.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
  Root: olympic, Root Metric: 10, Metric: 20
    track-item: glacier.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
3 nodes

```


show isis database

Syntax	<pre>show isis database <system-id> <brief detail extensive> <instance <i>instance-name</i>> <level (1 2)> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis database <system-id> <brief detail extensive> <level (1 2)> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display the entries in the IS-IS link-state database, which contains data about PDU packets.
Options	<p>none—Display standard information about IS-IS link-state database entries for all routing instances.</p> <p><i>system id</i>—(Optional) Display IS-IS link-state database entries for the specified intermediate system.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display IS-IS link-state database entries for the specified routing instance.</p> <p>level (1 2)—(Optional) Display IS-IS link-state database entries for the specified IS-IS level.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display standard information about IS-IS link-state database entries for all logical systems or for a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear isis database on page 436
List of Sample Output	<p>show isis database on page 533</p> <p>show isis database brief on page 534</p> <p>show isis database detail on page 534</p> <p>show isis database extensive on page 534</p>

Output Fields Table 43 on page 532 describes the output fields for the **show isis database** command. Output fields are listed in the approximate order in which they appear. Fields that contain internal IS-IS information useful only in troubleshooting obscure problems are not described in the table. For more details about these fields, contact your customer support representative.

Table 43: show isis database Output Fields

Field Name	Field Description	Level of Output
Interface name	Name of the interface on which the link-state PDU has been received; always IS-IS for this command.	All levels
level	Level of intermediate system: <ul style="list-style-type: none"> • 1—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system. • 2—Intermediate system routes between areas and toward other ASs. 	All levels
LSP ID	Link-state PDU identifier.	All levels
Sequence	Sequence number of the link-state PDU.	All levels
Checksum	Checksum value of the link-state PDU.	All levels
Lifetime (secs)	Remaining lifetime of the link-state PDU, in seconds.	All levels
Attributes	Attributes of the specified database: L1 , L2 , Overload , or Attached (L1 only).	none brief
# LSPs	Total number of link-state PDUs in the specified link-state database.	none brief
IP prefix	Prefix advertised by this link-state PDU.	detail extensive
IS neighbor	IS-IS neighbor of the advertising system.	detail extensive
ES neighbor	(J Series routers only) An ES-IS neighbor of the advertising system.	detail extensive
IP prefix	IPv4 prefix advertised by this link-state PDU.	detail extensive
V6 prefix	IPv6 prefix advertised by this link-state PDU.	detail extensive
Metric	Metric of the prefix or neighbor.	detail extensive
Header	<ul style="list-style-type: none"> • LSP ID—Link state PDU identifier of the header. • Length—Header length. • Allocated Length—Amount of length available for the header. • Router ID—Address of the local routing device. • Remaining Lifetime—Remaining lifetime of the link-state PDU, in seconds. 	extensive

Table 43: show isis database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Packet	<ul style="list-style-type: none"> • LSP ID—The identifier for the link-state PDU. • Length—Packet length. • Lifetime—Remaining lifetime, in seconds. • Checksum—The checksum of the link-state PDU. • Sequence—The sequence number of the link-state PDU. Every time the link-state PDU is updated, this number increments. • Attributes—Packet attributes. • NLPID—Network layer protocol identifier. • Fixed length—Specifies the set length for the packet. 	extensive
TLVs	<ul style="list-style-type: none"> • Area Address—Area addresses that the routing device can reach. • Speaks—Supported routing protocols. • IP router id—ID of the routing device (usually the IP address). • IP address—IPv4 address. • Hostname—Assigned name of the routing device. • IP prefix—IP prefix of the routing device. • Metric—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device. • IP extended prefix—Extended IP prefix of the routing device. • IS neighbor—Directly attached neighbor's name and metric. • IS extended neighbor—Directly attached neighbor's name, metric, and IP address. 	extensive

Sample Output

show isis database

```

user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00            0x3     0x3167     1057 L1 L2
camaro.00-00           0x5     0x770e     1091 L1 L2
ranier.00-00           0x4     0xaa95     1091 L1 L2
glacier.00-00          0x4     0x206f     1089 L1 L2
glacier.02-00          0x1     0xd141     1089 L1 L2
badlands.00-00        0x3     0x87a2     1093 L1 L2
  6 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00            0x6     0x8d6b     1096 L1 L2
camaro.00-00           0x9     0x877b     1101 L1 L2
ranier.00-00           0x8     0x855d     1103 L1 L2
glacier.00-00          0x7     0xf892     1098 L1 L2
glacier.02-00          0x1     0xd141     1089 L1 L2
badlands.00-00        0x6     0x562      1105 L1 L2
  6 LSPs

```

show isis database brief

The output for the **show isis database brief** command is identical to that for the **show isis database** command. For sample output, see [show isis database on page 533](#).

show isis database detail

```
user@host> show isis database logical-system CE3 sisira.00-00 detail
```

IS-IS level 1 link-state database:

```
sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 975 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  ES neighbor: 0015.0015.0015           Metric:      10 Down
  ES neighbor: 0025.0025.0025           Metric:      10 Down
  ES neighbor: 0030.0030.0030           Metric:      10 Down
  ES neighbor: 0040.0040.0040           Metric:      10 Down
  ES neighbor: sisira                     Metric:       0
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
```

IS-IS level 2 link-state database:

```
sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 993 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 50.50.50.50/32              Metric:      10 Internal Up
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
                                          Metric:       0 Internal Up
```

show isis database extensive

```
user@host> show isis database logical-system CE3 sisira.00-00 extensive
```

IS-IS level 1 link-state database:

```
sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 970 secs
```

```

IS neighbor: hemantha-CE3.02          Metric:      10
Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
ES neighbor: 0015.0015.0015          Metric:      10 Down
ES neighbor: 0025.0025.0025          Metric:      10 Down
ES neighbor: 0030.0030.0030          Metric:      10 Down
ES neighbor: 0040.0040.0040          Metric:      10 Down
ES neighbor: sisira                   Metric:       0
IP prefix: 1.0.0.0/24                 Metric:     10 External Down
IP prefix: 3.0.0.0/24                 Metric:     10 External Down
IP prefix: 4.0.0.0/24                 Metric:     10 External Down
IP prefix: 5.0.0.0/24                 Metric:     10 Internal Up
IP prefix: 15.15.15.15/32             Metric:     10 External Down
IP prefix: 25.25.25.25/32             Metric:     10 External Down
IP prefix: 30.30.30.30/32             Metric:     10 External Down
IP prefix: 40.40.40.40/32             Metric:     10 External Down
IP prefix: 60.60.60.60/32             Metric:       0 Internal Up

```

```

Header: LSP ID: sisira.00-00, Length: 336 bytes
Allocated length: 336 bytes, Router ID: 0.0.0.0
Remaining lifetime: 970 secs, Level: 1, Interface: 333
Estimated free bytes: 144, Actual free bytes: 0
Aging timer expires in: 970 secs
Protocols: IP, IPv6, CLNS

```

```

Packet: LSP ID: sisira.00-00, Length: 336 bytes, Lifetime : 1198 secs
Checksum: 0x10fc, Sequence: 0x11, Attributes: 0xb L1 L2 Attached
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 18, Packet version: 1, Max area: 0

```

```

TLVs:
Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
Speaks: IP
Speaks: IPV6
Speaks: CLNP
Hostname: sisira
ES neighbor TLV: Internal, Metric: default 0, Up
  ES: sisira
IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
IS extended neighbor: hemantha-CE3.02, Metric: default 10
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0040.0040.0040
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0025.0025.0025
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0015.0015.0015
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0030.0030.0030
IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down
IP extended prefix: 3.0.0.0/24 metric 10 down
IP extended prefix: 40.40.40.40/32 metric 10 down
IP extended prefix: 4.0.0.0/24 metric 10 down
IP extended prefix: 25.25.25.25/32 metric 10 down
IP extended prefix: 15.15.15.15/32 metric 10 down
IP extended prefix: 1.0.0.0/24 metric 10 down

```

```

IP extended prefix: 30.30.30.30/32 metric 10 down
IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
IP extended prefix: 60.60.60.60/32 metric 0 up
IP extended prefix: 5.0.0.0/24 metric 10 up
No queued transmissions

```

IS-IS level 2 link-state database:

```

sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 988 secs
IS neighbor: hemantha-CE3.02 Metric: 10
Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
IP prefix: 1.0.0.0/24 Metric: 10 External Down
IP prefix: 3.0.0.0/24 Metric: 10 External Down
IP prefix: 4.0.0.0/24 Metric: 10 External Down
IP prefix: 5.0.0.0/24 Metric: 10 Internal Up
IP prefix: 15.15.15.15/32 Metric: 10 External Down
IP prefix: 25.25.25.25/32 Metric: 10 External Down
IP prefix: 30.30.30.30/32 Metric: 10 External Down
IP prefix: 40.40.40.40/32 Metric: 10 External Down
IP prefix: 50.50.50.50/32 Metric: 10 Internal Up
IP prefix: 60.60.60.60/32 Metric: 0 Internal Up
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
Metric: 0 Internal Up

```

```

Header: LSP ID: sisira.00-00, Length: 427 bytes
Allocated length: 427 bytes, Router ID: 0.0.0.0
Remaining lifetime: 988 secs, Level: 2, Interface: 333
Estimated free bytes: 130, Actual free bytes: 0
Aging timer expires in: 988 secs
Protocols: IP, IPv6, CLNS

```

```

Packet: LSP ID: sisira.00-00, Length: 427 bytes, Lifetime : 1198 secs
Checksum: 0x69ac, Sequence: 0x13, Attributes: 0x3 L1 L2
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 20, Packet version: 1, Max area: 0

```

TLVs:

```

Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
Speaks: IP
Speaks: IPV6
Speaks: CLNP
Hostname: sisira
IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
IS extended neighbor: hemantha-CE3.02, Metric: default 10
IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
IP external prefix: 1.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down

```

```
IP extended prefix: 3.0.0.0/24 metric 10 down
IP extended prefix: 40.40.40.40/32 metric 10 down
IP extended prefix: 4.0.0.0/24 metric 10 down
IP extended prefix: 25.25.25.25/32 metric 10 down
IP extended prefix: 15.15.15.15/32 metric 10 down
IP extended prefix: 1.0.0.0/24 metric 10 down
IP extended prefix: 30.30.30.30/32 metric 10 down
ISO prefix-neighbor TLV: Internal, Metric: default 0, Up
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
IP prefix: 50.50.50.50/32, Internal, Metric: default 10, Up
IP extended prefix: 60.60.60.60/32 metric 0 up
IP extended prefix: 5.0.0.0/24 metric 10 up
IP extended prefix: 50.50.50.50/32 metric 10 up
No queued transmissions
```

show isis hostname

Syntax	show isis hostname <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis hostname
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display IS-IS hostname database information.
Options	<p>none—Display IS-IS hostname database information.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show isis hostname on page 538
Output Fields	Table 44 on page 538 describes the output fields for the show isis hostname command. Output fields are listed in the approximate order in which they appear.

Table 44: show isis hostname Output Fields

Field Name	Field Description
System Id	System identifier mapped to the hostname.
Hostname	Hostname mapped to the system identifier.
Type	Type of mapping between system identifier and hostname. <ul style="list-style-type: none"> Dynamic—Hostname mapping determined as described in RFC 2763, <i>Dynamic Hostname Exchange Mechanism for IS-IS</i>. Static—Hostname mapping configured by user.

Sample Output

show isis hostname

```

user@host> show isis hostname
IS-IS hostname database:
System Id      Hostname                                     Type
1921.6800.4201 isis1                         Dynamic
1921.6800.4202 isis2                         Static
1921.6800.4203 isis3                         Dynamic

```


show isis interface


Syntax	show isis interface <brief detail extensive> <interface-name> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis interface <brief detail extensive> <interface-name>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display status information about IS-IS-enabled interfaces.
	<div>  <p>NOTE: If the configured metric for an IS-IS level is above 63, and the wide-metrics-only statement is not configured, the show isis interface detail command and the show isis interface extensive command display 63 as the metric value for that level. Configure the wide-metrics-only statement to generate metric values greater than 63 on a per IS-IS level basis.</p> <p>The show isis interface command displays the configured metric value for an IS-IS level irrespective of whether is configured or not.</p> </div>
Options	<p>none—Display standard information about all IS-IS-enabled interfaces.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>interface-name—(Optional) Display information about the specified interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Example: Enabling Wide IS-IS Metrics for Traffic Engineering
List of Sample Output	show isis interface on page 541 show isis interface brief on page 541 show isis interface detail on page 542 show isis interface extensive on page 542 show isis interface extensive (With LDP) on page 542
Output Fields	Table 45 on page 540 describes the output fields for the show isis interface command. Output fields are listed in the approximate order in which they appear.

Table 45: show isis interface Output Fields

Field Name	Field Description	Level of Output
<i>interface-name</i>	Name of the interface.	detail
Designated router	Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.	detail
Index	Interface index assigned by the Junos OS kernel.	detail
State	Internal implementation information.	detail
Circuit id	Circuit identifier.	detail
Circuit type	Circuit type: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 	detail
LSP interval	Interval between link-state PDUs sent from the interface.	detail
CSNP interval	Interval between complete sequence number PDUs sent from the interface.	detail extensive
Sysid	System identifier.	detail
Interface	Interface through which the adjacency is made.	none brief
L or Level	Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 	All levels
CirID	Circuit identifier.	none brief
Level 1 DR	Level 1 designated intermediate system.	none brief
Level 2 DR	Level 2 designated intermediate system.	none brief
L1/L2 Metric	Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0.	none brief
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their label-switched paths (LSPs).	detail extensive
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs.	detail extensive
Adjacencies	Number of adjacencies established on this interface.	detail

Table 45: show isis interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Priority	Priority value for this interface.	detail
Metric	Metric value for this interface.	detail
Hello(s) / Hello Interval	Interface's hello interval.	detail extensive
Hold(s) / Hold Time	Interface's hold time.	detail extensive
Designated Router	Router responsible for sending network link-state advertisements, which describe all the routers attached to the network.	detail
Hello padding	Type of hello padding: <ul style="list-style-type: none"> • Adaptive—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state TLV. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. • Loose—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. • Strict—Padding is performed on all interface types and for all adjacency states, and is continuous. 	extensive
LDP sync state	Current LDP synchronization state: in sync , in holddown , or not supported .	extensive
reason	Reason for being in the LDP sync state.	extensive
config holdtime	Configured value of the hold timer.	extensive
remaining	If the state is not in sync and the hold time is not infinity, then this field displays the remaining hold time in seconds.	extensive

Sample Output

show isis interface

```

user@host> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
at-2/3/0.0         3   0x1 Point to Point    Point to Point    10/10
1o0.0              0   0x1 Passive           Passive           0/0

```

show isis interface brief

The output for the **show isis interface brief** command is identical to that for the **show isis interface** command. For sample output, see [show isis interface on page 541](#).

show isis interface detail

```
user@host> show isis interface detail
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1           1           64      10      9.000      27
    2           1           64      10      9.000      27
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1           0           64      0      0 Passive
    2           0           64      0      0 Passive
```

show isis interface extensive

```
user@host> show isis interface extensive
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s, Loose Hello padding
  Level 1
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled, Loose Hello padding
  Level 1
    Adjacencies: 0, Priority: 64, Metric: 0
    Passive
  Level 2
    Adjacencies: 0, Priority: 64, Metric: 0
    Passive
```

show isis interface extensive (With LDP)

```
user@host> show isis interface extensive
IS-IS interface database:
so-1/1/2.0
  Index: 114, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 20 s, Loose Hello padding
  Adjacency advertisement: Advertise
  LDP sync state: in sync, for: 00:01:28, reason: LDP up during config
  config holdtime: 20 seconds
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 11
    Hello Interval: 9.000 s, Hold Time: 27 s
    IPV4 MulticastMetric: 10
    IPV6 UnicastMetric: 10
```

show isis overview

Syntax	show isis overview <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis overview <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 8.5. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display IS-IS overview information.
Options	none —Display standard overview information about IS-IS for all routing instances. instance <i>instance-name</i> —(Optional) Display overview information for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show isis overview on page 544
Output Fields	Table 46 on page 543 lists the output fields for the show isis overview command. Output fields are listed in the approximate order in which they appear.

Table 46: show isis overview Output Fields

Field Name	Field Description
Instance	IS-IS routing instance.
Router ID	Router ID of the routing device.
Adjacency holddown	Adjacency holddown capability: enabled or disabled .
Maximum Areas	Maximum number of IS-IS areas advertised by the routing device.
LSP life time	Lifetime of the link-state PDU, in seconds.
Attached bit evaluation	Attached bit capability: enabled or disabled .
SPF delay	Delay before performing consecutive shortest-path-first (SPF) calculations.
SPF holddown	Delay before performing additional SPF calculations after the maximum number of consecutive SPF calculations is reached.

Table 46: show isis overview Output Fields (*continued*)

Field Name	Field Description
SPF rapid runs	Maximum number of SPF calculations that can be performed in succession before the holddown timer begins.
Overload bit at startup is set	Overload bit capability is enabled.
Overload high metrics	Overload high metrics capability: enabled or disabled .
Overload timeout	Time period after which overload is reset and the time that remains before the timer is set to expire.
Traffic engineering	Traffic engineering capability: enabled or disabled .
Restart	Graceful restart capability: enabled or disabled .
Restart duration	Time period for complete reacquisition of IS-IS neighbors.
Helper mode	Graceful restart helper capability: enabled or disabled .
Level	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 information • 2—Level 2 information
IPv4 is enabled	IP Protocol version 4 capability is enabled.
IPv6 is enabled	IP Protocol version 6 capability is enabled.
CLNS is enabled	(J Series routers only) OSI CLNP capability is enabled.
Internal route preference	Preference value of internal routes.
External route preference	Preference value of external routes.
Wide area metrics are enabled	Wide area metrics capability is enabled.
Narrow metrics are enabled	Narrow metrics capability is enabled.

Sample Output

show isis overview

```

user@host> show isis overview
Instance: master
Router ID: 192.168.1.220
Adjacency holddown: enabled

```

```
Maximum Areas: 3
LSP life time: 65535
Attached bit evaluation: enabled
SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
Overload bit at startup is set
  Overload high metrics: disabled
    Overload timeout: 300 sec, expires in 295 seconds
IPv4 is enabled, IPv6 is enabled
Traffic engineering: enabled
Restart: Enabled
  Restart duration: 210 sec
  Helper mode: Enabled
Level 1
  Internal route preference: 15
  External route preference: 160
  Wide metrics are enabled, Narrow metrics are enabled
Level 2
  Internal route preference: 18
  External route preference: 165
  Wide metrics are enabled
```

show isis route

Syntax	<pre>show isis route <destination> <inet inet6> <instance instance-name> <logical-system (all logical-system-name)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis route <destination> <inet inet6> <instance instance-name> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display the routes in the IS-IS routing table.
Options	<p>none—Display all routes in the IS-IS routing table for all supported address families for all routing instances.</p> <p>destination—(Optional) Destination address for the route.</p> <p>inet inet6—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.</p> <p>instance instance-name—(Optional) Display routes for the specified routing instance only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.</p>
Required Privilege Level	view
List of Sample Output	<p>show isis route logical-system on page 547</p> <p>show isis route (CLNS) on page 547</p> <p>show isis route on page 548</p>
Output Fields	<p>Table 47 on page 546 describes the output fields for the show isis route command. Output fields are listed in the approximate order in which they appear.</p>

Table 47: show isis route Output Fields

Field Name	Field Description
Current version	Number of the current version of the IS-IS routing table.

Table 47: show isis route Output Fields (*continued*)

Field Name	Field Description
L1	Version of Level 1 SPF that was run.
L2	Version of Level 2 SPF that was run.
Prefix	Destination of the route.
L	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2
Version	Version of SPF that generated the route.
Metric	Metric value associated with the route.
Type	Metric type: int (internal) or ext (external).
Interface	Interface to the next hop.
Via	System identifier of the next hop, displayed as a name if possible.
ISO Routes	ISO routing table entries.
snpa	MAC address.

Sample Output

show isis route logical-system

```

user@host> show isis route logical-system ls1
IS-IS routing table           Current version: L1: 8 L2: 11
Prefix      L Version Metric Type Interface  Via
10.9.7.0/30  2      11    20 int  gr-0/2/0.0  h
10.9.201.1/32 2      11    60 int  gr-0/2/0.0  h
IPV6 Unicast IS-IS routing table   Current version: L1: 9 L2: 11
Prefix      L Version Metric Type Interface  Via
8009:3::a09:3200/126 2      11    20 int  gr-0/2/0.0  h

```

show isis route (CLNS)

```

user@host> show isis route
IS-IS routing table           Current version: L1: 10 L2: 8
IPv4/IPv6 Routes
Prefix      L Version Metric Type Interface  Via
0.0.0.0/0   1      10    10 int  fe-0/0/1.0  ISIS.0
ISO Routes
Prefix L   Version Metric Type Interface  Via  snpa
0/0    1      10    10 int  fe-0/0/1.0  isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001/104

```

```

1          10          0 int
47.0005.80ff.f800.0000.0108.0001.1921.6800.4001/152
1          10          10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001.1921.6800.4002/152
1          10          20 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0002/104
1          10          0 int
47.0005.80ff.f800.0000.0108.0002.1921.6800.4001/152
1          10          10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56

```

show isis route

user@host> show isis route

```

IS-IS routing table          Current version: L1: 4 L2: 13
IPv4/IPv6 Routes
-----
Prefix                      L   Version  Metric Type Interface      NH   Via
10.255.71.52/32             2    13        10   int  ae0.0                 IPV4 camaro
10.255.71.238/32           2    13        20   int  so-6/0/0.0           IPV4 olympic
                               as0.0                 IPV4 glacier
10.255.71.239/32           2    13        20   int  so-6/0/0.0           IPV4 olympic
                               ae0.0                 IPV4 camaro
10.255.71.242/32           2    13        10   int  as0.0                 IPV4 glacier
10.255.71.243/32           2    13        10   int  so-6/0/0.0           IPV4 olympic
12.13.0.0/30                2    13        20   int  so-6/0/0.0           IPV4 olympic
12.15.0.0/30                2    13        20   int  so-6/0/0.0           IPV4 olympic
13.15.0.0/30                2    13        30   int  ae0.0                 IPV4 camaro
                               so-6/0/0.0           IPV4 olympic
                               as0.0                 IPV4 glacier
13.16.0.0/30                2    13        25   int  as0.0                 IPV4 glacier
14.15.0.0/30                2    13        20   int  ae0.0                 IPV4 camaro
192.2.1.0/30                2    13        30   int  so-6/0/0.0           IPV4 olympic
                               as0.0                 IPV4 glacier
1eee::/64                   2    13        30   int  so-6/0/0.0           IPV6 olympic
                               as0.0                 IPV6 glacier
abcd::10:255:71:52/128     2    13        10   int  ae0.0                 IPV6 camaro
abcd::10:255:71:238/128    2    13        20   int  so-6/0/0.0           IPV6 olympic
                               as0.0                 IPV6 glacier
abcd::10:255:71:239/128    2    13        20   int  so-6/0/0.0           IPV6 olympic

```

					ae0.0	IPV6 camaro
abcd::10:255:71:242/128	2	13	10	int	as0.0	IPV6 glacier
abcd::10:255:71:243/128	2	13	10	int	so-6/0/0.0	IPV6 olympic

show isis spf

Syntax	show isis spf (brief log results) <instance <i>instance-name</i> > <level (1 2)> <logical-system (all <i>logical-system-name</i>)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)>
Syntax (EX Series Switches)	show isis spf (brief log results) <instance <i>instance-name</i> > <level (1 2)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about IS-IS shortest-path-first (SPF) calculations.
Options	<p>brief—Display an overview of SPF calculations.</p> <p>instance <i>instance instance-name</i>—(Optional) Display SPF calculations for the specified routing instance.</p> <p>level (1 2)—(Optional) Display SPF calculations for the specified IS-IS level.</p> <p>log—Display the log of SPF calculations.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>results—Display the results of SPF calculations.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display SPF calculations for the specified topology only.</p>
Required Privilege Level	view
List of Sample Output	show isis spf log on page 551 show isis spf results logical-system on page 552 show isis spf results (CLNS) on page 553
Output Fields	Table 48 on page 550 describes the output fields for the show isis spf command. Output fields are listed in the approximate order in which they appear.

Table 48: show isis spf Output Fields

Field Name	Field Description
Node	System ID of a node.
Metric	Metric to the node.

Table 48: show isis spf Output Fields (*continued*)

Field Name	Field Description
Interface	Interface of the next hop.
Via	System ID of the next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
Start time	(log option only) Time that the SPF computation started.
Elapsed (secs)	(log option only) Length of time, in seconds, required to complete the SPF computation.
Count	(log option only) Number of times the SPF was triggered.
Reason	(log option only) Reason that the SPF computation was completed.

Sample Output

show isis spf log

```

user@host> show isis spf log logical-system lsl
IS-IS level 1 SPF log:
Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18 0.000069    1 Reconfig
Fri Oct 31 12:41:18 0.000107    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18 0.000050    3 Address change on so-1/2/2.0
Fri Oct 31 12:41:23 0.000033    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28 0.000178    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59 0.000060    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30 0.000161    2 Multi area attachment change
Fri Oct 31 12:56:58 0.000198    1 Periodic SPF
Fri Oct 31 13:10:29 0.000209    1 Periodic SPF
IS-IS level 2 SPF log:
Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18 0.000035    1 Reconfig
Fri Oct 31 12:41:18 0.000047    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18 0.000043    5 Address change on gr-0/2/0.0
Fri Oct 31 12:41:23 0.000022    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59 0.000144    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30 0.000257    3 New LSP skag.00-00
Fri Oct 31 12:54:37 0.000195    1 Periodic SPF
Fri Oct 31 12:55:50 0.000178    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55 0.000174    1 Updated LSP h.00-00
Fri Oct 31 12:55:58 0.000176    1 Updated LSP skag.00-00
Fri Oct 31 13:08:14 0.000198    1 Periodic SPF
IPv6 Unicast IS-IS level 1 SPF log:
Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18 0.000028    1 Reconfig
Fri Oct 31 12:41:18 0.000043    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18 0.000112    4 Updated LSP fix.00-00
Fri Oct 31 12:41:23 0.000059    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25 0.000041    1 Updated LSP fix.00-00

```

```

Fri Oct 31 12:41:28      0.000103    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59      0.000040    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30      0.000118    2 Multi area attachment change
Fri Oct 31 12:56:08      0.000289    1 Periodic SPF
Fri Oct 31 13:11:07      0.000214    1 Periodic SPF
IPV6 Unicast IS-IS level 2 SPF log:

```

```

Start time      Elapsed (secs) Count Reason
Fri Oct 31 12:41:18      0.000027    1 Reconfig
Fri Oct 31 12:41:18      0.000039    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18      0.000049    6 Updated LSP fix.00-00
Fri Oct 31 12:41:23      0.000025    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25      0.000023    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59      0.000087    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30      0.000123    3 New LSP skag.00-00
Fri Oct 31 12:55:50      0.000121    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55      0.000121    1 Updated LSP h.00-00
Fri Oct 31 12:55:58      0.000121    1 Updated LSP skag.00-00
Fri Oct 31 13:09:46      0.000201    1 Periodic SPF
...

```

show isis spf results logical-system

```
user@host> show isis spf results logical-system ls1
```

IS-IS level 1 SPF results:

Node	Metric	Interface	Via	SNPA
scat.00	10	ge-1/1/0.0	scat	0:90:69:a6:48:9d
	20	10.9.1.0/30		
fix.02	10			
fix.00	0			
	10	10.9.1.0/30		
	10	10.9.5.0/30		
	10	10.9.6.0/30		
	20	10.9.7.0/30		
	60	10.9.201.1/32		

3 nodes

IS-IS level 2 SPF results:

Node	Metric	Interface	Via	SNPA
skag.00	20	gr-0/2/0.0	h	
	30	10.9.7.0/30		
skag.02	20	gr-0/2/0.0	h	
h.00	10	gr-0/2/0.0	h	
	20	10.9.6.0/30		
	20	10.9.7.0/30		
	60	10.9.201.1/32		
fix.00	0			
	10	10.9.1.0/30		
	10	10.9.5.0/30		
	10	10.9.6.0/30		

4 nodes

IPV6 Unicast IS-IS level 1 SPF results:

Node	Metric	Interface	Via	SNPA
scat.00	10	ge-1/1/0.0	scat	0:90:69:a6:48:9d
		ge-1/1/0.0	scat	0:90:69:a6:48:9d
	20	8009:1::a09:1400/126		
fix.02	10			
fix.00	0			
	10	8009:1::a09:1400/126		
	10	8009:2::a09:1e00/126		

```

                20      8009:3::a09:3200/126
                10      8009:4::a09:2800/126
    3 nodes

IPv6 Unicast IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00    20      gr-0/2/0.0    h
           30      8009:3::a09:3200/126
skag.02    20      gr-0/2/0.0    h
           20      gr-0/2/0.0    h
h.00       10      gr-0/2/0.0    h
           20      8009:3::a09:3200/126
           20      8009:4::a09:2800/126
fix.00     0
           10      8009:1::a09:1400/126
           10      8009:2::a09:1e00/126
           10      8009:4::a09:2800/126
    4 nodes

Multicast IS-IS level 1 SPF results:
Node      Metric      Interface      Via      SNPA
scat.00    10      ge-1/1/0.0    scat    0:90:69:a6:48:9d
fix.02     10
fix.00     0
    3 nodes

Multicast IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00    20      gr-0/2/0.0    h
skag.02    20      gr-0/2/0.0    h
h.00       10      gr-0/2/0.0    h
fix.00     0
    4 nodes
...

```

show isis spf results (CLNS)

```

user@host> show isis spf results
IS-IS level 1 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00 10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           20      192.168.37.64/29
           10      1921.6800.4001
           20      1921.6800.4002
pro1-a.02 10
pro1-a.00 0
           0      10.255.245.1/32
           10      192.168.37.64/29
           0      1921.6800.4211
    3 nodes

IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00 10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           20      10.255.245.1/32
           20      192.168.37.64/29
           20      47.0005.80ff.f800.0000.0109.0010/104

```

pro1-a.02	10	
pro1-a.00	0	
	0	10.255.245.1/32
	10	192.168.37.64/29
3 nodes		

show isis statistics

Syntax	show isis statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display statistics about IS-IS traffic.
Options	<p>none—Display IS-IS traffic statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear isis statistics on page 440
List of Sample Output	show isis statistics on page 557
Output Fields	Table 49 on page 556 describes the output fields for the show isis statistics command. Output fields are listed in the approximate order in which they appear.

Table 49: show isis statistics Output Fields

Field Name	Field Description
PDU type	<p>PDU type:</p> <ul style="list-style-type: none"> • CSNP—Complete sequence number PDUs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU. • IIH—IS-IS hello packets are broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems. • LSP—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area. • PSNP—Partial sequence number PDUs are sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device. • Unknown—The PDU type is unknown.
Received	Number of PDUs received since IS-IS started or since the statistics were set to zero.
Processed	Number of PDUs received less the number dropped.
Drops	Number of PDUs dropped.
Sent	Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.
Rexmit	Number of PDUs retransmitted since IS-IS started or since the statistics were set to zero.
Total packets received/sent	Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.
SNP queue length	Number of CSPN and PSNP packets currently waiting in the queue for processing. This value is almost always 0.
LSP queue length	Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.
SPF runs	Number of shortest-path-first (SPF) calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable.
Fragments rebuilt	Number of link-state PDU fragments that the local system has computed.
LSP regenerations	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software decides that a link-state PDU must be removed from the network.

Sample Output

show isis statistics

```
user@host> show isis statistics
```

```
IS-IS statistics for merino:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	12227	12227	0	8184	683
IIH	113808	113808	0	115817	0
CSNP	198868	198868	0	198934	0
PSNP	6985	6979	6	8274	0
Unknown	0	0	0	0	0
Totals	331888	331882	6	331209	683

```
Total packets received: 331888 Sent: 331892
```

```
SNP queue length:      0 Drops:      0  
LSP queue length:      0 Drops:      0
```

```
SPF runs:              1014  
Fragments rebuilt:     1038  
LSP regenerations:     425  
Purges initiated:      0
```

show ospf3 database

Syntax	<pre>show ospf3 database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <external> <instance instance-name> <inter-area-prefix> <inter-area-router> <intra-area-prefix> <link> <link-local> <logical-system (all logical-system-name)> <lsa-id lsa-id> <network> <nssa> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <router></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show ospf3 database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <external> <instance instance-name> <inter-area-prefix> <inter-area-router> <intra-area-prefix> <link> <link-local> <lsa-id lsa-id> <network> <nssa> <router></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>advertising-router (address self) option introduced in Junos Release 9.5.</p> <p>advertising-router (address self) option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the OSPF version 3 (OSPFv3) link-state database, which contains data about link-state advertisement (LSA) packets.
Options	<p>none—Display standard information about all entries in the OSPFv3 link-state database.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>advertising-router (address self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.</p>

area *area-id*—(Optional) Display the LSAs in a particular area.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

inter-area-prefix—(Optional) Display information about interarea-prefix LSAs.

inter-area-router—(Optional) Display information about interarea-router LSAs.

intra-area-prefix—(Optional) Display information about intra-area-prefix LSAs.

link—(Optional) Display information about link LSAs.

link-local—(Optional) Display information about link-local LSAs.

logical-system (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(Optional) Display information about the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family other than IPv6 unicast, which is the default.

router—(Optional) Display information about router LSAs.

Required Privilege Level view

Related Documentation • [clear \(ospf | ospf3\) database on page 420](#)

List of Sample Output [show ospf3 database brief on page 564](#)
[show ospf3 database extensive on page 564](#)
[show ospf3 database summary on page 567](#)

Output Fields [Table 50 on page 559](#) lists the output fields for the **show ospf3 database** command. Output fields are listed in the approximate order in which they appear.

Table 50: show ospf3 database Output Fields

Field Name	Field Description	Level of Output
OSPF link state database, area <i>area-number</i>	Entries in the link-state database for this area.	brief detail extensive
OSPF AS SCOPE link state database	Entries in the AS scope link-state database.	brief detail extensive

Table 50: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
OSPF Link-Local link state database, interface <i>interface-name</i>	Entries in the link-local link-state database for this interface.	brief detail extensive
area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Type	Type of link advertisement: Extern , InterArPfx , InterArRtr , IntraArPrx , Link , Network , NSSA , or Router .	brief detail extensive
ID	Link identifier included in the advertisement. An asterisk (*) preceding the identifier marks database entries that originated from the local routing device.	brief detail extensive
Adv Rtr	Address of the routing device that sent the advertisement.	brief detail extensive
Seq	Link sequence number of the advertisement.	brief detail extensive
Age	Time elapsed since the LSA was originated, in seconds.	brief detail extensive
Cksum	Checksum value of the LSA.	brief detail extensive
Len	Length of the advertisement, in bytes.	brief detail extensive
Router (Router Link-State Advertisements)		
bits	Flags describing the routing device that generated the LSP.	detail extensive
Options	Option bits carried in the router LSA.	detail extensive
For Each Router Link		
Type	Type of interface. The value of all other output fields describing a routing device interface depends on the interface's type: <ul style="list-style-type: none"> • PointToPoint (1)—Point-to-point connection to another routing device. • Transit (2)—Connection to a transit network. • Virtual (4)—Virtual link. 	detail extensive
Loc-if-id	Local interface ID assigned to the interface that uniquely identifies the interface with the routing device.	detail extensive
Nbr-if-id	Interface ID of the neighbor's interface for this routing device link.	detail extensive
Nbr-rtr-id	Router ID of the neighbor routing device (for type 2 interfaces, the attached link's designated router).	detail extensive
Metric	Cost of the router link.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive

Table 50: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
Network (Network Link-State Advertisements)		
Options	Option bits carried in the network LSA.	detail extensive
Attached Router	Router IDs of each of the routing devices attached to the link. Only routing devices that are fully adjacent to the designated router are listed. The designated router includes itself in this list.	detail extensive
InterArPfx (Interarea-Prefix Link-State Advertisements)		
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
InterArRtr (Interarea-Router Link-State Advertisements)		
Dest-router-id	Router ID of the routing device described by the LSA.	detail extensive
options	Optional capabilities supported by the routing device.	detail extensive

Table 50: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	detail extensive
Prefix	IPv6 address prefix.	extensive
Prefix-options	Option bit associated with the prefix.	extensive
Extern (External Link-State Advertisements)		
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of the route, which depends on the value of Type .	detail extensive
Type <i>n</i>	Type of external metric: Type 1 or Type 2 .	detail extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Link (Link-State Advertisements)		
IPv6-Address	IPv6 link-local address on the link for which this link LSA originated.	detail extensive
Options	Option bits carried in the link LSA.	detail extensive
priority	Router priority of the interface attaching the originating routing device to the link.	detail extensive
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	detail extensive
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive

Table 50: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
IntraArPfx (Intra-Area-Prefix Link-State Advertisements)		
Ref-lsa-type	LSA type of the referenced LSA. <ul style="list-style-type: none"> Router—Address prefixes are associated with a router LSA. Network—Address prefixes are associated with a network LSA. 	detail extensive
Ref-lsa-id	Link-state ID of the referenced LSA.	detail extensive
Ref-router-id	Advertising router ID of the referenced LSA.	detail extensive
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	detail extensive
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of this prefix. Expressed in the same units as the interface costs in the router LSAs.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>hh:mm:ss</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>hh:mm:ss</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>hh:mm:ss</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
<i>n</i> Router LSAs	Number of router LSAs in the link-state database.	summary
<i>n</i> Network LSAs	Number of network LSAs in the link-state database.	summary
<i>n</i> InterArPfx LSAs	Number of interarea-prefix LSAs in the link-state database.	summary

Table 50: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>n</i> InterArRtr LSAs	Number of interarea-router LSAs in the link-state database.	summary
<i>n</i> IntraArPfx LSAs	Number of intra-area-prefix LSAs in the link-state database.	summary
Externals	Display of the external LSA database.	summary
<i>n</i> Extern LSAs	Number of external LSAs in the link-state database.	summary
Interface <i>interface-name</i>	Name of the interface for which link-local LSA information is displayed.	summary
<i>n</i> Link LSAs	Number of link LSAs in the link-state database.	summary

Sample Output

show ospf3 database brief

```

user@host> show ospf3 database brief
      OSPF3 link state database, area 0.0.0.0
      Type      ID          Adv Rtr      Seq          Age  Cksum  Len
      Router    0.0.0.1      10.255.4.85  0x80000003   885  0xa697  40
      Router    *0.0.0.1     10.255.4.93  0x80000002   953  0xc677  40
      InterArPfx *0.0.0.2     10.255.4.93  0x80000001   910  0xb96f  44
      InterArRtr *0.0.0.1     10.255.4.93  0x80000001   910  0xe159  32
      IntraArPfx *0.0.0.1     10.255.4.93  0x80000002   432  0x788f  72

      OSPF3 link state database, area 0.0.0.1
      Type      ID          Adv Rtr      Seq          Age  Cksum  Len
      Router    *0.0.0.1     10.255.4.93  0x80000003   916  0xea40  40
      Router    0.0.0.1     10.255.4.97  0x80000006   851  0xc95b  40
      Network    0.0.0.2     10.255.4.97  0x80000002   916  0x4598  32
      InterArPfx *0.0.0.1     10.255.4.93  0x80000002   117  0xa980  44
      InterArPfx *0.0.0.2     10.255.4.93  0x80000002    62  0xd47e  44
      NSSA      0.0.0.1     10.255.4.97  0x80000002   362  0x45ee  44
      IntraArPfx 0.0.0.1     10.255.4.97  0x80000006   851  0x2f77  52

      OSPF3 AS SCOPE link state database
      Type      ID          Adv Rtr      Seq          Age  Cksum  Len
      Extern    0.0.0.1     10.255.4.85  0x80000002    63  0x9b86  44
      Extern    *0.0.0.1     10.255.4.93  0x80000001   910  0x59c9  44

      OSPF3 Link-Local link state database, interface ge-1/3/0.0
      Type      ID          Adv Rtr      Seq          Age  Cksum  Len
      Link      *0.0.0.2     10.255.4.93  0x80000003   916  0x4dab  64

```

show ospf3 database extensive

```

user@host> show ospf3 database extensive
      OSPF3 link state database, area 0.0.0.0
      Type      ID          Adv Rtr      Seq          Age  Cksum  Len
      Router    0.0.0.1     10.255.4.85  0x80000003  1028  0xa697  40
      bits 0x2, Options 0x13
      Type PointToPoint (1), Metric 10

```

```

    Loc-If-Id 2, Nbr-If-Id 3, Nbr-Rtr-Id 10.255.4.93
    Aging timer 00:42:51
    Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
Router    *0.0.0.1          10.255.4.93      0x80000002  1096  0xc677  40
    bits 0x3, Options 0x13
    Type PointToPoint (1), Metric 10
    Loc-If-Id 3, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.85
    Gen timer 00:00:40
    Aging timer 00:41:44
    Installed 00:18:16 ago, expires in 00:41:44, sent 00:18:14 ago
    Ours
InterArPfx *0.0.0.2          10.255.4.93      0x80000001  1053  0xb96f  44
    Prefix feee::10:10:2:0/126
    Prefix-options 0x0, Metric 10
    Gen timer 00:17:02
    Aging timer 00:42:26
    Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
    Ours
InterArPfx *0.0.0.3          10.255.4.93      0x80000001  1053  0x71d3  44
    Prefix feee::10:255:4:97/128
    Prefix-options 0x0, Metric 10
    Gen timer 00:21:07
    Aging timer 00:42:26
    Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
    Ours
InterArRtr *0.0.0.1          10.255.4.93      0x80000001  1053  0xe159  32
    Dest-router-id 10.255.4.97, Options 0x19, Metric 10
    Gen timer 00:29:18
    Aging timer 00:42:26
    Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
    Ours
IntraArPfx 0.0.0.1          10.255.4.85      0x80000002  1028  0x2403  72
    Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.85
    Prefix-count 2
    Prefix feee::10:255:4:85/128
    Prefix-options 0x2, Metric 0
    Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
    Aging timer 00:42:51
    Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
IntraArPfx *0.0.0.1          10.255.4.93      0x80000002   575  0x788f  72
    Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.93
    Prefix-count 2
    Prefix feee::10:255:4:93/128
    Prefix-options 0x2, Metric 0
    Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
    Gen timer 00:33:23
    Aging timer 00:50:24
    Installed 00:09:35 ago, expires in 00:50:25, sent 00:09:33 ago
    OSPF3 link state database, area 0.0.0.1
    Type      ID          Adv Rtr          Seq          Age  Cksum  Len
Router    *0.0.0.1          10.255.4.93      0x80000003  1059  0xea40  40
    bits 0x3, Options 0x19
    Type Transit (2), Metric 10
    Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
    Gen timer 00:08:51
    Aging timer 00:42:20
    Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Router      0.0.0.1          10.255.4.97      0x80000006   994  0xc95b  40
    bits 0x2, Options 0x19

```

```

Type Transit (2), Metric 10
  Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
Network 0.0.0.2 10.255.4.97 0x80000002 1059 0x4598 32
Options 0x11
  Attached router 10.255.4.97
  Attached router 10.255.4.93
Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
InterArPfx *0.0.0.1 10.255.4.93 0x80000002 260 0xa980 44
  Prefix feee::10:10:1:0/126
  Prefix-options 0x0, Metric 10
  Gen timer 00:45:39
  Aging timer 00:55:39
  Installed 00:04:20 ago, expires in 00:55:40, sent 00:04:18 ago
  Ours
InterArPfx *0.0.0.2 10.255.4.93 0x80000002 205 0xd47e 44
  Prefix feee::10:255:4:93/128
  Prefix-options 0x0, Metric 0
  Gen timer 00:46:35
  Aging timer 00:56:35
  Installed 00:03:25 ago, expires in 00:56:35, sent 00:03:23 ago
  Ours
InterArPfx *0.0.0.3 10.255.4.93 0x80000001 1089 0x9bbb 44
  Prefix feee::10:255:4:85/128
  Prefix-options 0x0, Metric 10
  Gen timer 00:04:46
  Aging timer 00:41:51
  Installed 00:18:09 ago, expires in 00:41:51, sent 00:17:43 ago
  Ours
NSSA 0.0.0.1 10.255.4.97 0x80000002 505 0x45ee 44
  Prefix feee::200:200:1:0/124
  Prefix-options 0x8, Metric 10, Type 2,
  Aging timer 00:51:35
  Installed 00:08:22 ago, expires in 00:51:35, sent 02:37:54 ago
IntraArPfx 0.0.0.1 10.255.4.97 0x80000006 994 0x2f77 52
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:255:4:97/128
  Prefix-options 0x2, Metric 0
  Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
IntraArPfx 0.0.0.3 10.255.4.97 0x80000002 1059 0x4446 52
  Ref-lsa-type Network, Ref-lsa-id 0.0.0.2, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:10:2:0/126
  Prefix-options 0x0, Metric 0
  Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
  OSPF3 AS SCOPE link state database
  Type ID Adv Rtr Seq Age Cksum Len
Extern 0.0.0.1 10.255.4.85 0x80000002 206 0x9b86 44
  Prefix feee::100:100:1:0/124
  Prefix-options 0x0, Metric 20, Type 2,
  Aging timer 00:56:34
  Installed 00:03:23 ago, expires in 00:56:34, sent 02:37:54 ago
Extern *0.0.0.1 10.255.4.93 0x80000001 1053 0x59c9 44
  Prefix feee::200:200:1:0/124
  Prefix-options 0x0, Metric 10, Type 2,
  Gen timer 00:25:12

```

```

Aging timer 00:42:26
Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago

  OSPF3 Link-Local link state database, interface ge-1/3/0.0
Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link      *0.0.0.2      10.255.4.93  0x80000003   1059 0x4dab 64
fe80::290:69ff:fe39:1cdb
Options 0x11, priority 128
Prefix-count 1
Prefix feee::10:10:2:0/126 Prefix-options 0x0
Gen timer 00:12:56
Aging timer 00:42:20
Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Link      0.0.0.2      10.255.4.97  0x80000003   205  0xa87d 64
fe80::290:69ff:fe38:883e
Options 0x11, priority 128
Prefix-count 1
Prefix feee::10:10:2:0/126 Prefix-options 0x0
Aging timer 00:56:35
Installed 00:03:22 ago, expires in 00:56:35, sent 02:37:54 ago

  OSPF3 Link-Local link state database, interface so-2/2/0.0
Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link      0.0.0.2      10.255.4.85  0x80000002   506  0x42bb 64
fe80::280:42ff:fe10:f169
Options 0x13, priority 128
Prefix-count 1
Prefix feee::10:10:1:0/126 Prefix-options 0x0
Aging timer 00:51:34
Installed 00:08:23 ago, expires in 00:51:34, sent 02:37:54 ago
Link      *0.0.0.3      10.255.4.93  0x80000002   505  0x6b7a 64
fe80::280:42ff:fe10:f177
Options 0x13, priority 128
Prefix-count 1
Prefix feee::10:10:1:0/126 Prefix-options 0x0
Gen timer 00:37:28
Aging timer 00:51:35
Installed 00:08:25 ago, expires in 00:51:35, sent 00:08:23 ago
Ours

```

show ospf3 database summary

```

user@host> show ospf3 database summary
Area 0.0.0.0:
  2 Router LSAs
  1 InterArPfx LSAs
  1 InterArRtr LSAs
  1 IntraArPfx LSAs
Area 0.0.0.1:
  2 Router LSAs
  1 Network LSAs
  2 InterArPfx LSAs
  1 NSSA LSAs
  1 IntraArPfx LSAs
Externals:
  2 Extern LSAs
Interface ge-1/3/0.0:
  1 Link LSAs
Interface lo0.0:

```

Interface so-2/2/0.0:
1 Link LSAs

show ospf database

Syntax	<pre>show ospf database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <asbrsummary> <external> <instance instance-name> <link-local> <logical-system (all logical-system-name)> <lsa-id lsa-id> <netsummary> <network> <nssa> <opaque-area> <router></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show ospf database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <asbrsummary> <external> <instance instance-name> <link-local> <lsa-id lsa-id> <netsummary> <network> <nssa> <opaque-area> <router></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>advertising-router self (address self) option introduced in Junos OS Release 9.5.</p> <p>advertising-router self (address self) option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the OSPF version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets.
Options	<p>none—Display standard information about entries in the OSPFv2 link-state database for all routing instances.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>advertising-router (address self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.</p> <p>area area-id—(Optional) Display the LSAs in a particular area.</p>

asbrsummary—(Optional) Display summary AS boundary router LSA entries.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

link-local—(Optional) Display information about link-local LSAs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

netsummary—(Optional) Display summary network LSAs.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Display opaque area-scope LSAs.

router—(Optional) Display information about router LSAs.

Required Privilege Level

view

Related Documentation

- [clear \(ospf | ospf3\) database on page 420](#)

List of Sample Output

[show ospf database on page 572](#)
[show ospf database brief on page 572](#)
[show ospf database detail on page 572](#)
[show ospf database extensive on page 574](#)
[show ospf database summary on page 576](#)

Output Fields

[Table 51 on page 570](#) describes the output fields for the **show ospf database** command. Output fields are listed in the approximate order in which they appear.

Table 51: show ospf database Output Fields

Field Name	Field Description	Level of Output
area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Type	Type of link advertisement: ASBRSum , Extern , Network , NSSA , OpaqArea , Router , or Summary .	All levels
ID	LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device.	All levels
Adv Rtr	Address of the routing device that sent the advertisement.	All levels
Seq	Link sequence number of the advertisement.	All levels

Table 51: show ospf database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Age	Time elapsed since the LSA was originated, in seconds.	All levels
Opt	Optional OSPF capabilities associated with the LSA.	All levels
Cksum	Checksum value of the LSA.	All levels
Len	Length of the advertisement, in bytes.	All levels
Router	Router link-state advertisement information: <ul style="list-style-type: none"> • bits—Flags describing the routing device that generated the LSP. • link count—Number of links in the advertisement. • id—ID of a routing device or subnet on the link. • data—For stub networks, the subnet mask. Otherwise, the IP address of the routing device that generated the LSP. • type—Type of link. It can be PointToPoint, Transit, Stub, or Virtual. • TOS count—Number of type-of-service (ToS) entries in the advertisement. • TOS 0 metric—Metric for ToS 0. • TOS—Type-of-service (ToS) value. • metric—Metric for the ToS. 	detail extensive
Network	Network link-state advertisement information: <ul style="list-style-type: none"> • mask—Network mask. • attached router—ID of the attached neighbor. 	detail extensive
Summary	Summary link-state advertisement information: <ul style="list-style-type: none"> • mask—Network mask. • TOS—Type-of-service (ToS) value. • metric—Metric for the ToS. 	detail extensive
Gen timer	How long until the LSA is regenerated.	extensive
Aging timer	How long until the LSA expires.	extensive
Installed <i>hh:mm:ss</i> ago	How long ago the route was installed.	extensive
expires in <i>hh:mm:ss</i>	How long until the route expires.	extensive
sent <i>hh:mm:ss</i> ago	How long ago the LSA was sent.	extensive
Last changed <i>hh:mm:ss</i> ago	How long ago the route was changed.	extensive
Change count	Number of times the route has changed.	extensive

Table 51: show ospf database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Ours	Indicates that this is a local advertisement.	extensive
Router LSAs	Number of router link-state advertisements in the link-state database.	summary
Network LSAs	Number of network link-state advertisements in the link-state database.	summary
Summary LSAs	Number of summary link-state advertisements in the link-state database.	summary
NSSA LSAs	Number of not-so-stubby area link-state advertisements in the link-state database.	summary

Sample Output

show ospf database

```

user@host> show ospf database
OSPF link state database, Area 0.0.0.1
  Type      ID            Adv Rtr      Seq          Age    Opt  Cksum  Len
Router     10.255.70.103   10.255.70.103 0x80000002   215    0x20 0x4112  48
Router     *10.255.71.242  10.255.71.242 0x80000002   214    0x20 0x11b1  48
Summary    *23.1.1.0       10.255.71.242 0x80000002   172    0x20 0x6d72  28
Summary    *24.1.1.0       10.255.71.242 0x80000002   177    0x20 0x607e  28
NSSA       *33.1.1.1       10.255.71.242 0x80000002   217    0x28 0x73bd  36

      OSPF link state database, Area 0.0.0.2
  Type      ID            Adv Rtr      Seq          Age    Opt  Cksum  Len
Router     10.255.71.52   10.255.71.52  0x80000004   174    0x20 0xd021  36
Router     *10.255.71.242  10.255.71.242 0x80000003   173    0x20 0xe191  36
Network    *23.1.1.1       10.255.71.242 0x80000002   173    0x20 0x9c76  32
Summary    *12.1.1.0       10.255.71.242 0x80000001   217    0x20 0xfeec  28
Summary    *24.1.1.0       10.255.71.242 0x80000002   177    0x20 0x607e  28
NSSA       *33.1.1.1       10.255.71.242 0x80000001   222    0x28 0xe047  36

      OSPF link state database, Area 0.0.0.3
  Type      ID            Adv Rtr      Seq          Age    Opt  Cksum  Len
Router     10.255.71.238   10.255.71.238 0x80000003   179    0x20 0x3942  36
Router     *10.255.71.242  10.255.71.242 0x80000003   177    0x20 0xf37d  36
Network    *24.1.1.1       10.255.71.242 0x80000002   177    0x20 0xc591  32
Summary    *12.1.1.0       10.255.71.242 0x80000001   217    0x20 0xfeec  28
Summary    *23.1.1.0       10.255.71.242 0x80000002   172    0x20 0x6d72  28
NSSA       *33.1.1.1       10.255.71.242 0x80000001   222    0x28 0xeb3b  36

```

show ospf database brief

The output for the **show ospf database brief** command is identical to that for the **show ospf database** command. For sample output, see [show ospf database on page 572](#).

show ospf database detail

```

user@host> show ospf database detail
OSPF link state database, Area 0.0.0.1
  Type      ID            Adv Rtr      Seq          Age    Opt  Cksum  Len
Router     10.255.70.103   10.255.70.103 0x80000002   261    0x20 0x4112  48

```

```

bits 0x0, link count 2
id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
TOS count 0, TOS 0 metric 1
id 12.1.1.0, data 255.255.255.0, Type Stub (3)
TOS count 0, TOS 0 metric 1
Router *10.255.71.242 10.255.71.242 0x80000002 260 0x20 0x11b1 48
bits 0x3, link count 2
id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
TOS count 0, TOS 0 metric 1
id 12.1.1.0, data 255.255.255.0, Type Stub (3)
TOS count 0, TOS 0 metric 1
Summary *23.1.1.0 10.255.71.242 0x80000002 218 0x20 0x6d72 28
mask 255.255.255.0
TOS 0x0, metric 1
Summary *24.1.1.0 10.255.71.242 0x80000002 223 0x20 0x607e 28
mask 255.255.255.0
TOS 0x0, metric 1
NSSA *33.1.1.1 10.255.71.242 0x80000002 263 0x28 0x73bd 36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0

```

OSPF link state database, Area 0.0.0.2

Type	ID	Adv Rtr	Seq	Age	Opt	Cksum	Len
Router	10.255.71.52	10.255.71.52	0x80000004	220	0x20	0xd021	36
bits 0x0, link count 1							
id 23.1.1.1, data 23.1.1.2, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Router	*10.255.71.242	10.255.71.242	0x80000003	219	0x20	0xe191	36
bits 0x3, link count 1							
id 23.1.1.1, data 23.1.1.1, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Network	*23.1.1.1	10.255.71.242	0x80000002	219	0x20	0x9c76	32
mask 255.255.255.0							
attached router 10.255.71.242							
attached router 10.255.71.52							
Summary	*12.1.1.0	10.255.71.242	0x80000001	263	0x20	0xfeec	28
mask 255.255.255.0							
TOS 0x0, metric 1							
Summary	*24.1.1.0	10.255.71.242	0x80000002	223	0x20	0x607e	28
mask 255.255.255.0							
TOS 0x0, metric 1							
NSSA	*33.1.1.1	10.255.71.242	0x80000001	268	0x28	0xe047	36
mask 255.255.255.255							
Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0							

OSPF link state database, Area 0.0.0.3

Type	ID	Adv Rtr	Seq	Age	Opt	Cksum	Len
Router	10.255.71.238	10.255.71.238	0x80000003	225	0x20	0x3942	36
bits 0x0, link count 1							
id 24.1.1.1, data 24.1.1.2, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Router	*10.255.71.242	10.255.71.242	0x80000003	223	0x20	0xf37d	36
bits 0x3, link count 1							
id 24.1.1.1, data 24.1.1.1, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Network	*24.1.1.1	10.255.71.242	0x80000002	223	0x20	0xc591	32
mask 255.255.255.0							
attached router 10.255.71.242							
attached router 10.255.71.238							
Summary	*12.1.1.0	10.255.71.242	0x80000001	263	0x20	0xfeec	28
mask 255.255.255.0							

```

TOS 0x0, metric 1
Summary *23.1.1.0      10.255.71.242    0x80000002    218  0x20 0x6d72  28
mask 255.255.255.0
TOS 0x0, metric 1
NSSA  *33.1.1.1      10.255.71.242    0x80000001    268  0x28 0xeb3b  36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0

```

show ospf database extensive

```

user@host> show ospf database extensive
  OSPF link state database, Area 0.0.0.1
  Type      ID          Adv Rtr          Seq      Age  Opt  Cksum  Len
Router 10.255.70.103  10.255.70.103    0x80000002    286  0x20 0x4112  48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:14
  Installed 00:04:43 ago, expires in 00:55:14
  Last changed 00:04:43 ago, Change count: 2
Router *10.255.71.242  10.255.71.242    0x80000002    285  0x20 0x11b1  48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
  Gen timer 00:45:15
  Aging timer 00:55:15
  Installed 00:04:45 ago, expires in 00:55:15, sent 00:04:43 ago
  Last changed 00:04:45 ago, Change count: 2, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243  0x20 0x6d72  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0      10.255.71.242    0x80000002    248  0x20 0x607e  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA  *33.1.1.1      10.255.71.242    0x80000002    288  0x28 0x73bd  36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:48 ago
Last changed 00:04:48 ago, Change count: 2, Ours

  OSPF link state database, Area 0.0.0.2
  Type      ID          Adv Rtr          Seq      Age  Opt  Cksum  Len
Router 10.255.71.52   10.255.71.52     0x80000004    245  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:55

```

```

    Installed 00:04:02 ago, expires in 00:55:55
    Last changed 00:04:02 ago, Change count: 2
Router *10.255.71.242    10.255.71.242    0x80000003    244    0x20 0xe191    36
    bits 0x3, link count 1
    id 23.1.1.1, data 23.1.1.1, Type Transit (2)
    TOS count 0, TOS 0 metric 1
    Gen timer 00:45:56
    Aging timer 00:55:56
    Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
    Last changed 00:04:04 ago, Change count: 2, Ours
Network *23.1.1.1      10.255.71.242    0x80000002    244    0x20 0x9c76    32
    mask 255.255.255.0
    attached router 10.255.71.242
    attached router 10.255.71.52
    Gen timer 00:45:56
    Aging timer 00:55:56
    Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
    Last changed 00:04:04 ago, Change count: 1, Ours
Summary *12.1.1.0      10.255.71.242    0x80000001    288    0x20 0xfeec    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:12
    Aging timer 00:55:12
    Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:04 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0      10.255.71.242    0x80000002    248    0x20 0x607e    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:52
    Aging timer 00:55:52
    Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:04 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1        10.255.71.242    0x80000001    293    0x28 0xe047    36
    mask 255.255.255.255
    Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0
    Gen timer 00:45:07
    Aging timer 00:55:07
    Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:04 ago
    Last changed 00:04:53 ago, Change count: 1, Ours

    OSPF link state database, Area 0.0.0.3
    Type      ID      Adv Rtr      Seq      Age  Opt  Cksum  Len
Router  10.255.71.238    10.255.71.238    0x80000003    250  0x20 0x3942    36
    bits 0x0, link count 1
    id 24.1.1.1, data 24.1.1.2, Type Transit (2)
    TOS count 0, TOS 0 metric 1
    Aging timer 00:55:50
    Installed 00:04:07 ago, expires in 00:55:50
    Last changed 00:04:07 ago, Change count: 2
Router *10.255.71.242    10.255.71.242    0x80000003    248  0x20 0xf37d    36
    bits 0x3, link count 1
    id 24.1.1.1, data 24.1.1.1, Type Transit (2)
    TOS count 0, TOS 0 metric 1
    Gen timer 00:45:52
    Aging timer 00:55:52
    Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
    Last changed 00:04:08 ago, Change count: 2, Ours
Network *24.1.1.1      10.255.71.242    0x80000002    248  0x20 0xc591    32
    mask 255.255.255.0
    attached router 10.255.71.242
    attached router 10.255.71.238

```

```
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 1, Ours
Summary *12.1.1.0      10.255.71.242    0x80000001    288  0x20 0xfeec  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:13 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243  0x20 0x6d72  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA  *33.1.1.1        10.255.71.242    0x80000001    293  0x28 0xeb3b  36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:13 ago
Last changed 00:04:53 ago, Change count: 1, Ours
```

show ospf database summary

```
user@host> show ospf database summary
Area 0.0.0.1:
  2 Router LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.2:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.3:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Externals:
Interface fe-2/2/1.0:
Interface ge-0/3/2.0:
Interface so-0/1/2.0:
Interface so-0/1/2.0:
```

show policy damping

Syntax	show policy damping <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show policy damping
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display information about BGP route flap damping parameters.
Options	none —Display information about BGP route flap damping parameters. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Additional Information	In the output from this command, figure-of-merit values correlate with the probability of future instability of a routing device. 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.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • “Configuring BGP Flap Damping Parameters” in the Routing Policy Configuration Guide • clear bgp damping on page 428 • show route damping on page 614
List of Sample Output	show policy damping on page 578
Output Fields	Table 52 on page 577 describes the output fields for the show policy damping command. Output fields are listed in the approximate order in which they appear.

Table 52: show policy damping Output Fields

Field Name	Field Description
Halflife	Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. 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.

Table 52: show policy damping Output Fields (*continued*)

Field Name	Field Description
Reuse merit	Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.
Suppress/cutoff merit	Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When 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.
Maximum suppress time	Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability.
Computed values	<ul style="list-style-type: none"> • Merit ceiling—Maximum merit that a flapping route can collect. • Maximum decay—Maximum decay half-life, in minutes.

Sample Output

show policy damping

```

user@host> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "standard-damping":
  Halflife: 10 minutes
  Reuse merit: 4000 Suppress/cutoff merit: 8000
  Maximum suppress time: 30 minutes
  Computed values:
    Merit ceiling: 32120
    Maximum decay: 12453

```


show rip general-statistics

Syntax	show rip general-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show rip general-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display brief RIP statistics.
Options	<p>none—Display brief RIP statistics.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear rip general-statistics on page 443
List of Sample Output	show rip general-statistics on page 579
Output Fields	Table 53 on page 579 lists the output fields for the show rip general-statistics command. Output fields are listed in the approximate order in which they appear.

Table 53: show rip general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no rcv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIP.
max memory	Most memory used by RIP.

Sample Output

show rip general-statistics

```
user@host> show rip general-statistics
RIPv2 I/O info:
    bad msgs      :      0
    no rcv intf   :      0
```

```
curr memory : 0
max memory  : 0
```

show rip neighbor

Syntax	show rip neighbor <instance (all <i>instance-name</i>)> <logical-system (all <i>logical-system-name</i>)> < <i>name</i> >
Syntax (EX Series Switches and QFX Series)	show rip neighbor <instance (all <i>instance-name</i>)> < <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about RIP neighbors.
Options	<p>none—Display information about all RIP neighbors for all instances.</p> <p>instance (all <i>instance-name</i>)—(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>name</i>—(Optional) Display detailed information about only the specified RIP neighbor.</p>
Required Privilege Level	view
List of Sample Output	show rip neighbor on page 582 show rip neighbor (With Demand Circuits Configured) on page 582
Output Fields	Table 54 on page 581 lists the output fields for the show rip neighbor command. Output fields are listed in the approximate order in which they appear.

Table 54: show rip neighbor Output Fields

Field Name	Field Description
Neighbor	<p>Name of the RIP neighbor.</p> <p>NOTE: Beginning with Junos OS Release 11.1, when you configure demand circuits, the output displays a demand circuit (DC) flag next to neighbor interfaces configured for demand circuits.</p> <p>If you configure demand circuits at the [edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>] hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the [edit protocols rip group <i>group-name</i>] hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits.</p>

Table 54: show rip neighbor Output Fields (*continued*)

Field Name	Field Description
State	State of the connection: Up or Dn (Down).
Source Address	Address of the port on the local router.
Destination Address	Address of the port on the remote router.
Send Mode	Send options: broadcast , multicast , none , or version 1 .
Receive Mode	Type of packets to accept: both , none , version 1 , or version 2 .
In Met	Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.

Sample Output

show rip neighbor

```

user@host> show rip neighbor
Neighbor      Local  Source      Destination  Send  Receive  In
-----      -
ge-2/3/0.0    Up    192.168.9.105  192.168.9.107  bcast  both      1
at-5/1/1.42    Dn    (null)         (null)         mcast  v2 only   3
at-5/1/0.42    Dn    (null)         (null)         mcast  both      3
at-5/1/0.0     Up    20.0.0.1       224.0.0.9      mcast  both      3
so-0/0/0.0     Up    192.168.9.97   224.0.0.9      mcast  both      3

```

show rip neighbor (With Demand Circuits Configured)

```

user@host> show rip neighbor
Neighbor      Local  Source      Destination  Send  Receive  In
-----      -
so-0/1/0.0(DC) Up    10.10.10.2    224.0.0.9    mcast  both      1
so-0/2/0.0(DC) Up    13.13.13.2    224.0.0.9    mcast  both      1

```

show rip statistics

Syntax	<pre>show rip statistics <instance (all <i>instance-name</i>)> <logical-system (all <i>logical-system-name</i>)> <name> <peer (all <i>address</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show rip statistics <instance (all <i>instance-name</i>)> <name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display RIP statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.
Options	<p>none—Display RIP statistics for all routing instances.</p> <p>instance (all <i>instance-name</i>)—(Optional) Display RIP statistics for all instances or for only the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>name—(Optional) Display detailed information about only the specified RIP neighbor.</p> <p>peer (all <i>address</i>)—(Optional) Display RIP statistics for a single peer or all peers.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear rip statistics on page 444
List of Sample Output	show rip statistics on page 584
Output Fields	Table 55 on page 584 lists the output fields for the show rip statistics command. Output fields are listed in the approximate order in which they appear.

Table 55: show rip statistics Output Fields

Field Name	Field Description
RIP info	<p>Information about RIP on the specified interface:</p> <ul style="list-style-type: none"> • port—UDP port number used for RIP. • update interval—Interval between routing table updates, in seconds. • holddown—Hold-down interval, in seconds. • timeout—Timeout interval, in seconds. • restart in progress—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart. • restart time—Estimated time for the graceful restart to finish, in seconds. • restart will complete in—Remaining time for the graceful restart to finish, in seconds. • rts learned—Number of routes learned through RIP. • rts held down—Number of routes held down by RIP. • rqsts dropped—Number of received request packets that were dropped. • resps dropped—Number of received response packets that were dropped.
<i>logical-interface</i>	<p>Name of the logical interface and its statistics:</p> <ul style="list-style-type: none"> • routes learned—Number of routes learned on the logical interface. • routes advertised—Number of routes advertised by the logical interface.
Counter	<p>List of counter types:</p> <ul style="list-style-type: none"> • Updates Sent—Number of update messages sent. • Triggered Updates Sent—Number of triggered update messages sent. • Responses Sent—Number of response messages sent. • Bad Messages—Number of invalid messages received. • RIPv1 Updates Received—Number of RIPv1 update messages received. • RIPv1 Bad Route Entries—Number of RIPv1 invalid route entry messages received. • RIPv1 Updates Ignored—Number of RIPv1 update messages ignored. • RIPv2 Updates Received—Number of RIPv2 update messages received. • RIPv2 Bad Route Entries—Number of RIPv2 invalid route entry messages received. • RIPv2 Updates Ignored—Number of RIPv2 update messages ignored. • Authentication Failures—Number of received update messages that failed authentication. • RIP Requests Received—Number of RIP request messages received. • RIP Requests Ignored—Number of RIP request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

Sample Output

show rip statistics

```
user@host> show rip statistics so-0/0/0.0
```

```

RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s
restart in progress: restart time 60s; restart will complete in 55s
  rts learned  rts held down  rqsts dropped  resps dropped
            0            0            0            0

```

```

so-0/0/0.0: 0 routes learned; 501 routes advertised

```

Counter	Total	Last 5 min	Last minute
-----	-----	-----	-----
Updates Sent	0	0	0
Triggered Updates Sent	0	0	0
Responses Sent	0	0	0
Bad Messages	0	0	0
RIPv1 Updates Received	0	0	0
RIPv1 Bad Route Entries	0	0	0
RIPv1 Updates Ignored	0	0	0
RIPv2 Updates Received	0	0	0
RIPv2 Bad Route Entries	0	0	0
RIPv2 Updates Ignored	0	0	0
Authentication Failures	0	0	0
RIP Requests Received	0	0	0
RIP Requests Ignored	0	0	0

show ripng general-statistics

Syntax	show ripng general-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch)	show ripng general-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display general RIP next-generation (RIPng) statistics.
Options	none —Display general RIPng statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear ripng general-statistics on page 445
List of Sample Output	show ripng general-statistics on page 586
Output Fields	Table 56 on page 586 lists the output fields for the show ripng general-statistics command. Output fields are listed in the approximate order in which they appear.

Table 56: show ripng general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no recv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIPng.
max memory	Most memory used by RIPng.

Sample Output

show ripng general-statistics

```

user@host> show ripng general-statistics
RIPng I/O info:
  bad msgs      :      0
  no recv intf  :      0
  curr memory   :      0
  max memory    :      0

```


show ripng neighbor

Syntax	show ripng neighbor <logical-system (all <i>logical-system-name</i>)> < <i>name</i> >
Syntax (EX Series Switch)	show ripng neighbor < <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about RIP next-generation (RIPng) neighbors.
Options	<p>none—Display information about all RIPng neighbors.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>name</i>—(Optional) Display detailed information about a specific RIPng neighbor.</p>
Required Privilege Level	view
List of Sample Output	show ripng neighbor on page 588
Output Fields	Table 57 on page 587 lists the output fields for the show ripng neighbor command. Output fields are listed in the approximate order in which they appear.

Table 57: show ripng neighbor Output Fields

Field Name	Field Description
Neighbor	Name of RIPng neighbor.
State	State of the connection: Up or Dn (Down).
Source Address	Source address.
Destination Address	Destination address.
Send	Send options: broadcast , multicast , none , version 1 , or yes .
Recv	Type of packets to accept: both , none , version 1 , or yes .
In Met	Metric added to incoming routes when advertising into RIPng routes that were learned from other protocols.

Sample Output

show ripng neighbor

```
user@host> show ripng neighbor
```

Neighbor	State	Source Address	Dest Address	Send	Recv	In Met
-----	-----	-----	-----	-----	-----	-----
fe-0/0/2.0	Up	fe80::290:69ff:fe68:b002	ff02::9	yes	yes	1

show ripng statistics

Syntax	show ripng statistics <logical-system (all <i>logical-system-name</i>)> < <i>name</i> >
Syntax (EX Series Switch)	show ripng statistics < <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display RIP next generation (RIPng) statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.
Options	<p>none—Display RIPng statistics for all neighbors.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>name</i>—(Optional) Display detailed information about a specific RIPng neighbor.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ripng statistics on page 446
List of Sample Output	show ripng statistics on page 590
Output Fields	Table 58 on page 589 lists the output fields for the show ripng statistics command. Output fields are listed in the approximate order in which they appear.

Table 58: show ripng statistics Output Fields

Field Name	Field Description
RIPng info	<p>Information about RIPng on the specified interface:</p> <ul style="list-style-type: none"> • port—UDP port number used for RIPng. • holddown—Hold-down interval, in seconds. • rts learned—Number of routes learned through RIPng. • rts held down—Number of routes held down by RIPng. • rqsts dropped—Number of received request packets that were dropped. • resps dropped—Number of received response packets that were dropped. • restart—Graceful restart status. Displayed when RIPng is or has been in the process of graceful restart.

Table 58: show ripng statistics Output Fields (*continued*)

Field Name	Field Description
<i>logical-interface</i>	Name of the logical interface and its statistics: <ul style="list-style-type: none"> routes learned—Number of routes learned on the logical interface. routes advertised—Number of routes advertised by the logical interface. timeout—Timeout interval, in seconds. update interval—Interval between routing table updates, in seconds.
Counter	List of counter types: <ul style="list-style-type: none"> Updates Sent—Number of update messages sent. Triggered Updates Sent—Number of triggered update messages sent. Responses Sent—Number of response messages sent. Bad Messages—Number of invalid messages received. Updates Received—Number of RIPng update messages received. Bad Route Entries—Number of RIPng invalid route entry messages received. Updates Ignored—Number of RIPng update messages ignored. RIPng Requests Received—Number of RIPng request messages received. RIPng Requests Ignored—Number of RIPng request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

Sample Output

show ripng statistics

```

user@host> show ripng statistics
RIPng info: port 521; holddown 120s;
      rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0

so-0/1/3.0: 0 routes learned; 1 routes advertised; timeout 180s; update interval
20s
Counter              Total    Last 5 min  Last minute
-----
Updates Sent          934         16          4
Triggered Updates Sent    1          0          0
Responses Sent         0          0          0
Bad Messages           0          0          0
Updates Received        0          0          0
Bad Route Entries       0          0          0
Updates Ignored         0          0          0
RIPng Requests Received  0          0          0
RIPng Requests Ignored   0          0          0

```

show route

Syntax	<pre>show route <all> <destination-prefix> <logical-system (all logical-system-name)> <private></pre>
Syntax (EX Series Switches)	<pre>show route <all> <destination-prefix> <private></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option private introduced in Junos OS Release 9.5.</p> <p>Option private introduced in Junos OS Release 9.5 for EX Series switches.</p>
Description	Display the active entries in the routing tables.
Options	<p>none—Display brief information about all active entries in the routing tables.</p> <p>all—(Optional) Display information about all routing tables, including private, or internal, routing tables.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>private—(Optional) Display information only about all private, or internal, routing tables.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring RIP • Example: Configuring RIPng • Example: Configuring IS-IS • Examples: Configuring Internal BGP Peering • Examples: Configuring External BGP Peering • Examples: Configuring OSPF Routing Policy
List of Sample Output	<p>show route on page 594</p> <p>show route destination-prefix on page 594</p> <p>show route extensive on page 594</p>

Output Fields Table 59 on page 592 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 59: show route Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	<p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> • active (routes that are active). • holddown (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. <p>However, if you have configured advertisement of multiple routes (with the add-path or advertise-inactive statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong.</p> <ul style="list-style-type: none"> • hidden (routes that are not used because of a routing policy).
<i>destination-prefix</i>	<p>Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
<i>[protocol, preference]</i>	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>

Table 59: show route Output Fields (*continued*)

Field Name	Field Description
<i>weeks:days</i> <i>hours:minutes:seconds</i>	How long the route been known (for example, 2w4d 13:11:14 , or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
localpref	Local preference value included in the route.
from	Interface from which the route was received.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
to	<p>Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.</p> <p>If the destination is Discard, traffic is dropped.</p>

Table 59: show route Output Fields (*continued*)

Field Name	Field Description
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • lsp-path-name—Name of the LSP used to reach the next hop. • label-action—MPLS label and operation occurring at the next hop. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).

Sample Output

show route

```

user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
    * [MVPN/70] 19:53:41, metric2 1
    Indirect
1:65500:1:10.0.0.40/240
    * [BGP/170] 19:53:29, localpref 100, from 10.0.0.30
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
    * [BGP/170] 19:53:29, localpref 100, from 10.0.0.30
    AS path: I
    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
    [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
    AS path: I
    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF

```

show route destination-prefix

```

user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12    * [Static/5] 2w4d 12:54:27
                  > to 192.168.167.254 via fxp0.0

```

show route extensive

```

user@host> show route extensive

```



```

v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
    PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group 225.1.1.1

    Next hop type: Indirect
    Address: 0x92455b8
    Next-hop reference count: 2
    Source: 10.0.0.30
    Protocol next hop: 10.0.0.40
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
      Local AS: 65500 Peer AS: 65500
    Age: 3  Metric2: 1
    Task: BGP_65500.10.0.0.30+179
    Announcement bits (2): 0-PIM.v1 1-mvpn global task
    AS path: I (Originator) Cluster list: 10.0.0.30
    AS path: Originator ID: 10.0.0.40
    Communities: target:65520:100
    Import Accepted
    Localpref: 100
    Router ID: 10.0.0.30
    Primary Routing Table bgp.mvpn.0
    Indirect next hops: 1
      Protocol next hop: 10.0.0.40 Metric: 1
      Indirect next hop: 2 no-forward
      Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
      10.0.0.40/32 Originating RIB: inet.3
        Metric: 1                               Node path count: 1
        Forwarding nexthops: 1
          Nexthop: 10.0.24.4 via lt-0/3/0.24

```

show route active-path

Syntax	show route active-path <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route active-path <brief detail extensive terse>
Release Information	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.
Options	none —Display all active routes. brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route active-path on page 596 show route active-path brief on page 597 show route active-path detail on page 597 show route active-path extensive on page 598 show route active-path terse on page 600
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    * [Direct/0] 21:33:52
                  > via lo0.0
10.255.71.50/32    * [IS-IS/15] 00:18:13, metric 10
                  > to 100.1.2.1 via so-2/1/3.0
100.1.2.0/24      * [Direct/0] 00:18:36
                  > via so-2/1/3.0
100.1.2.2/32      * [Local/0] 00:18:41
                  Local via so-2/1/3.0
192.168.64.0/21   * [Direct/0] 21:33:52
```

```

> via fxp0.0
192.168.70.19/32  *Local/0] 21:33:52
                  Local via fxp0.0

```

show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see [show route active-path on page 596](#).

show route active-path detail

```

user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:37:10
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

    AS path: I

10.255.71.50/32 (1 entry, 1 announced)
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Router, Next hop index: 397
    Next-hop reference count: 4
    Next hop: 100.1.2.1 via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:31 Metric: 10
    Task: IS-IS
    Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
    AS path: I

100.1.2.0/24 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:54
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

    AS path: I

100.1.2.2/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 11
    Interface: so-2/1/3.0
    State: <Active NoReadvrt Int>

```

```
Local AS: 200
Age: 21:59
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I
```

show route active-path extensive

```
user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via lo0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {100.1.2.1}
IS-IS level 2, LSP fragment 0
*IS-IS Preference: 15
Level: 1
Next hop type: Router, Next hop index: 397
Next-hop reference count: 4
Next hop: 100.1.2.1 via so-2/1/3.0, selected
State: <Active Int>
```

```

Local AS: 200
Age: 24:08 Metric: 10
Task: IS-IS
Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
AS path: I

100.1.2.0/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:31
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
AS path: I

100.1.2.2/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 24:36
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

```

show route active-path terse

```
user@host> show route active-path terse
```

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.255.70.19/32	D	0			>1o0.0	
*	10.255.71.50/32	I	15	10		>100.1.2.1	
*	100.1.2.0/24	D	0			>so-2/1/3.0	
*	100.1.2.2/32	L	0			Local	
*	192.168.64.0/21	D	0			>fxp0.0	
*	192.168.70.19/32	L	0			Local	

show route all

Syntax	<code>show route all</code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route all</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about all routes in all routing tables, including private, or internal, tables.
Options	none —Display information about all routes in all routing tables, including private, or internal, tables. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route all on page 601
Output Fields	In Junos OS Release 9.5 and later, only the output fields for the show route all command display all routing tables, including private, or hidden, routing tables. The output field table of the show route command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later.

Sample Output

show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```
user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
1          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
2          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
800017     *[VPLS/7] 1d 14:00:16
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 14:00:26
            > via vt-3/2/0.32772, Pop

user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
```

```

+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
1          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
2          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
800017     *[VPLS/7] 1d 13:54:49
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 13:54:59
            > via vt-3/2/0.32772, Pop
vt-3/2/0.32769 [VPLS/7] 1d 13:54:49
              Unusable
vt-3/2/0.32772 [VPLS/7] 1d 13:54:59
              Unusable
    
```


show route aspath-regex

Syntax	<code>show route aspath-regex <i>regular-expression</i></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route aspath-regex <i>regular-expression</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that match the specified autonomous system (AS) path regular expression.
Options	<p><i>regular-expression</i>—Regular expression that matches an entire AS path.</p> <p><i>logical-system (all <i>logical-system-name</i>)</i>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	<p>You can specify a regular expression as:</p> <ul style="list-style-type: none"> • An individual AS number • A period wildcard used in place of an AS number • An AS path regular expression that is enclosed in parentheses <p>You also can include the operators described in the table of AS path regular expression operators in the <i>Junos Policy Framework Configuration Guide</i>. The following list summarizes these operators:</p> <ul style="list-style-type: none"> • <i>{m,n}</i>—At least <i>m</i> and at most <i>n</i> repetitions of the AS path term. • <i>{m}</i>—Exactly <i>m</i> repetitions of the AS path term. • <i>{m,}</i>—<i>m</i> or more repetitions of the AS path term. • <i>*</i>—Zero or more repetitions of an AS path term. • <i>+</i>—One or more repetitions of an AS path term. • <i>?</i>—Zero or one repetition of an AS path term. • <i>aspath_term aspath_term</i>—Match one of the two AS path terms. <p>When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:</p> <pre>show route aspath-regex ". * 234 . *"</pre>
Required Privilege Level	view

List of Sample Output [show route aspath-regex \(Matching a Specific AS Number\) on page 604](#)
 [show route aspath-regex \(Matching Any Path with Two AS Numbers\) on page 604](#)

Output Fields For information about output fields, see the output field table for the [show route](#) command.

Sample Output

[show route aspath-regex \(Matching a Specific AS Number\)](#)

```
user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25      *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25   *[IS-IS/15] 09:15:37, metric 37, tag 1
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
                   [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
...
```

[show route aspath-regex \(Matching Any Path with Two AS Numbers\)](#)

```
user@host> show route aspath-regex ?.* 234 3561.*?

inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 2685 2686 Incomplete
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 5696 7369 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 6327 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...
```

show route best

Syntax	<code>show route best <i>destination-prefix</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route best <i>destination-prefix</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . <i>destination-prefix</i> —Address or range of addresses. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route best on page 605 show route best detail on page 606 show route best extensive on page 607 show route best terse on page 607
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route best

```

user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[OSPF/10] 1d 13:19:20, metric 2
                  > to 10.31.1.6 via ge-3/1/0.0
                  via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[RSVP/7] 1d 13:20:13, metric 2
                  > via so-0/3/0.0, label-switched-path green-r1-r3

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

```

```
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8      *[Direct/0] 2d 01:43:34
                 > via fxp2.0
                 [Direct/0] 2d 01:43:34
                 > via fxp1.0
```

show route best detail

```
user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
    *OSPF    Preference: 10
              Next-hop reference count: 9
              Next hop: 10.31.1.6 via ge-3/1/0.0, selected
              Next hop: via so-0/3/0.0
              State: <Active Int>
              Local AS:    69
              Age: 1d 13:20:06      Metric: 2
              Area: 0.0.0.0
              Task: OSPF
              Announcement bits (2): 0-KRT 3-Resolve tree 2
              AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
    State: <FlashAll>
    *RSVP    Preference: 7
              Next-hop reference count: 5
              Next hop: via so-0/3/0.0 weight 0x1, selected
              Label-switched-path green-r1-r3
              Label operation: Push 100016
              State: <Active Int>
              Local AS:    69
              Age: 1d 13:20:59      Metric: 2
              Task: RSVP
              Announcement bits (1): 1-Resolve tree 2
              AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
    *Direct Preference: 0
              Next hop type: Interface
              Next-hop reference count: 1
              Next hop: via fxp2.0, selected
              State: <Active Int>
              Age: 2d 1:44:20
              Task: IF
              AS path: I
    Direct Preference: 0
              Next hop type: Interface
              Next-hop reference count: 1
              Next hop: via fxp1.0, selected
              State: <NotBest Int>
              Inactive reason: No difference
              Age: 2d 1:44:20
              Task: IF
              AS path: I
```

show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see [show route best detail on page 606](#).

show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.255.70.103/32  0 10      2          >10.31.1.6
                               so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.255.70.103/32  R  7      2          >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.0.0.0/8        D  0          >fxp2.0
                    D  0          >fxp1.0
```

show route brief

Syntax	<code>show route brief</code> <code><destination-prefix></code> <code><logical-system (all logical-system-name)></code>
Syntax (EX Series Switches)	<code>show route brief</code> <code><destination-prefix></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display brief information about the active entries in the routing tables.
Options	none —Display all active entries in the routing table. destination-prefix —(Optional) Display active entries for the specified address or range of addresses. logical-system (all logical-system-name) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route brief on page 608
Output Fields	For information about output fields, see the Output Field table of the show route command.

Sample Output

show route brief

```
user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                   Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                   > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.0.0/18     *[Static/5] 1w5d 20:30:29
                   > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                   > via fxp0.0
192.168.164.51/32  *[Local/0] 2w4d 13:11:14
                   Local via fxp0.0
207.17.136.192/32 *[Static/5] 2w4d 13:11:14
```

```

> to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16    *[Direct/0] 1w5d 20:30:28
                  > via fe-0/0/3.0
100.101.2.3/32   *[Local/0] 1w5d 20:30:28
                  Local via fe-0/0/3.0
224.0.0.5/32     *[OSPF/10] 1w5d 20:30:29, metric 1
                  MultiRecv
```

show route community

Syntax	<code>show route community <i>as-number:community-value</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route community <i>as-number:community-value</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community.
Options	<p><i>as-number:community-value</i>—One or more community identifiers. <i>as-number</i> is the AS number, and <i>community-value</i> is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks. Community identifiers can include wildcards.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	Specifying the community option displays all routes matching the community found within the routing table. The community option does not limit the output to only the routes being advertised to the neighbor after any egress routing policy.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show route detail on page 619
List of Sample Output	show route community on page 610
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route community

```
user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

4.0.0.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 1 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
6.0.0.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
```



```
9.2.0.0/16      AS Path: {666} 234 2548 568 721 Incomplete
                  to 192.156.169.1 via 192.156.169.14(so-0/0/0)
                  *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
                  AS Path: {666} 234 2548 1673 1675 1747 IGP
                  to 192.156.169.1 via 192.156.169.14(so-0/0/0)
```

show route community-name

Syntax	<code>show route community-name <i>community-name</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route community-name <i>community-name</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community, specified by a community name.
Options	<i>community-name</i> —Name of the community. brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route community-name on page 612
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route community-name

```
user@host> show route community-name red-com
inet.0: 17 destinations, 17 routes (16 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

instance1.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.212/32  *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                  AS path: 300 I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
20.20.20.20/32    *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                  AS path: I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
100.1.4.0/24     *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                  AS path: I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.204:10:10.255.245.212/32
    *[BGP/170] 00:06:40, localpref 100, from 10.255.245.204
        AS path: 300 I
        > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:20.20.20.20/32
    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
        AS path: I
        > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:100.1.4.0/24
    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
        AS path: I
        > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

instance1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route damping

Syntax	show route damping (decayed history suppressed) <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show route damping (decayed history suppressed) <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display the BGP routes for which updates might have been reduced because of route flap damping.
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>decayed—Display route damping entries that might no longer be valid, but are not suppressed.</p> <p>history—Display entries that have already been withdrawn, but have been logged.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>suppressed—Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear bgp damping on page 428 • show policy damping on page 577
List of Sample Output	show route damping decayed detail on page 617 show route damping history on page 618 show route damping history detail on page 618
Output Fields	Table 60 on page 614 lists the output fields for the show route damping command. Output fields are listed in the approximate order in which they appear.

Table 60: show route damping Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table—for example, inet.0 .	All levels
destinations	Number of destinations for which there are routes in the routing table.	All levels

Table 60: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holddown (routes that are in a pending state before being declared inactive) • hidden (the routes are not used because of a routing policy) 	All levels
<i>destination-prefix (entry, announced)</i>	Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination.	detail extensive
<i>[protocol, preference]</i>	Protocol from which the route was learned and the preference value for the route. <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>	All levels
Next-hop reference count	Number of references made to the next hop.	detail extensive
Source	IP address of the route source.	detail extensive
Next hop	Network layer address of the directly reachable neighboring system.	detail extensive
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	detail extensive
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.	detail extensive
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.	detail extensive
State	Flags for this route. For a description of possible values for this field, see the output field table for the show route detail command.	detail extensive
Local AS	AS number of the local routing device.	detail extensive
Peer AS	AS number of the peer routing device.	detail extensive

Table 60: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
Age	How long the route has been known.	detail extensive
Metric	Metric for the route.	detail extensive
Task	Name of the protocol that has added the route.	detail extensive
Announcement bits	List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only.	detail extensive
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> I—IGP. E—EGP. ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured. { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. ()—Parentheses enclose a confederation. ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>	All levels
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	brief none
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	brief none
Communities	Community path attribute for the route. See the output field table for the show route detail command.	detail extensive
Localpref	Local preference value included in the route.	All levels
Router ID	BGP router ID as advertised by the neighbor in the open message.	detail extensive
Merit (last update/now)	Last updated and current figure-of-merit value.	detail extensive

Table 60: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
damping-parameters	Name that identifies the damping parameters used, which is defined in the damping statement at the [edit policy-options] hierarchy level.	detail extensive
Last update	Time of most recent change in path attributes.	detail extensive
First update	Time of first change in path attributes, which started the route damping process.	detail extensive
Flaps	Number of times the route has gone up or down or its path attributes have changed.	detail extensive
Suppressed	(suppressed keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it.	All levels
Reusable in	(suppressed keyword only) Time when a suppressed route will again be available.	All levels
Preference will be	(suppressed keyword only) Preference value that will be applied to the route when it is again active.	All levels

Sample Output

show route damping decayed detail

```

user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
  *BGP      Preference: 170/-101
             Next-hop reference count: 151973
             Source: 172.23.2.129
             Next hop: via so-1/2/0.0
             Next hop: via so-5/1/0.0, selected
             Next hop: via so-6/0/0.0
             Protocol next hop: 172.23.2.129
             Indirect next hop: 89a1a00 264185
             State: <Active Ext>
             Local AS: 65000 Peer AS: 65490
             Age: 3:28      Metric2: 0
             Task: BGP_65490.172.23.2.129+179
             Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

6-Resolve tree 2 7-Resolve tree 3
AS path: 65490 65520 65525 65525 65525 I ()
Communities: 65501:390 65501:2000 65501:3000 65504:701
Localpref: 100
Router ID: 172.23.2.129
Merit (last update/now): 1934/1790
damping-parameters: damping-high
Last update:          00:03:28 First update:          00:06:40
Flaps: 2

```

show route damping history

```
user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15      [BGP ] 2d 22:47:58, localpref 100
                  AS path: 65220 65501 65502 I
                  > to 192.168.60.85 via so-3/1/0.0
```

show route damping history detail

```
user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
    BGP                /-101
        Next-hop reference count: 69058
        Source: 192.168.60.85
        Next hop: 192.168.60.85 via so-3/1/0.0, selected
        State: <Hidden Ext>
        Inactive reason: Unusable path
        Local AS: 65000 Peer AS: 65220
        Age: 2d 22:48:10
        Task: BGP_65220.192.168.60.85+179
        AS path: 65220 65501 65502 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:3561
        Localpref: 100
        Router ID: 192.168.80.25
        Merit (last update/now): 1000/932
        damping-parameters: set-normal
        Last update:          00:01:05 First update:          00:01:05
        Flaps: 1
```


show route detail

Syntax	show route detail <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route detail <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display detailed information about the active entries in the routing tables.
Options	<p>none—Display all active entries in the routing table on all systems.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route detail on page 627 show route detail (with BGP Multipath) on page 633
Output Fields	<p>Table 61 on page 619 describes the output fields for the show route detail command. Output fields are listed in the approximate order in which they appear.</p>

Table 61: show route detail Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active) holddown (routes that are in the pending state before being declared inactive) hidden (routes that are not used because of a routing policy)

Table 61: show route detail Output Fields (*continued*)

Field Name	Field Description
<i>route-destination</i> (entry, announced)	<p>Route destination (for example:10.0.0.1/24). The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[<i>protocol, preference</i>]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
Level	<p>(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 62 on page 623 .

Table 61: show route detail Output Fields (*continued*)

Field Name	Field Description
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path lsp-path-name	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 63 on page 625 .
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metricn	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.

Table 61: show route detail Output Fields (*continued*)

Field Name	Field Description
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table .
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • Recorded—The AS path is recorded by the sample process (sampled). • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • []—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 64 on page 627 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).

Table 61: show route detail Output Fields (*continued*)

Field Name	Field Description
control flags	Control flags: none or Site Down .
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 62 on page 623 describes all possible values for the **Next-hop Types** output field.

Table 62: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.

Table 62: Next-hop Types Output Field Values (*continued*)

Next-Hop Type	Description
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	<p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 63 on page 625 describes all possible values for the **State** output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 63: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).

Table 63: State Output Field Values (*continued*)

Value	Description
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 64 on page 627 describes the possible values for the **Communities** output field.

Table 64: Communities Output Field Values

Value	Description
<i>area-number</i>	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0 . A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
<i>bandwidth: local AS number:link-bandwidth-number</i>	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
<i>domain-id</i>	Unique configurable number that identifies the OSPF domain.
<i>domain-id-vendor</i>	Unique configurable number that further identifies the OSPF domain.
<i>link-bandwidth-number</i>	Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).
<i>local AS number</i>	Local AS number: from 1 through 65,535 .
<i>options</i>	1 byte. Currently this is only used if the route type is 5 or 7 . Setting the least significant bit in the field indicates that the route carries a type 2 metric.
<i>origin</i>	(Used with VPNs) Identifies where the route came from.
<i>ospf-route-type</i>	1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses.
<i>route-type-vendor</i>	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000 . The format is <i>area-number:ospf-route-type:options</i> .
<i>rte-type</i>	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306 . The format is <i>area-number:ospf-route-type:options</i> .
<i>target</i>	Defines which VPN the route participates in; target has the format <i>32-bit IP address:16-bit number</i> . For example, 10.19.0.0:100.
<i>unknown IANA</i>	Incoming IANA codes with a value between 0x1 and 0x7fff . This code of the BGP extended community attribute is accepted, but it is not recognized.
<i>unknown OSPF vendor community</i>	Incoming IANA codes with a value above 0x8000 . This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route detail

```
user@host> show route detail
```

```
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
```

```
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:31:43
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2
    Next hop: via so-0/3/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:30:17
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I
  OSPF Preference: 10
    Next-hop reference count: 1
    Next hop: via so-0/3/0.0, selected
    State: <Int>
    Inactive reason: Route Preference
    Local AS: 69
    Age: 1:30:17 Metric: 1
    Area: 0.0.0.0
    Task: OSPF
    AS path: I

10.31.1.1/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 7
    Interface: so-0/3/0.0
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:30:20
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
  *OSPF Preference: 10
    Next-hop reference count: 9
    Next hop: via so-0/3/0.0
    Next hop: 10.31.1.6 via ge-3/1/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:29:56 Metric: 2
    Area: 0.0.0.0
    Task: OSPF
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...
```

```

224.0.0.2/32 (1 entry, 1 announced)
  *PIM    Preference: 0
          Next-hop reference count: 18
          State: <Active NoReadvrt Int>
          Local AS:    69
          Age: 1:31:45
          Task: PIM Recv
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP   Preference: 0
          Next-hop reference count: 18
          State: <Active NoReadvrt Int>
          Local AS:    69
          Age: 1:31:43
          Task: IGMP
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
          Label-switched-path green-r1-r3
          Label operation: Push 100096
          State: <Active Int>
          Local AS:    69
          Age: 1:25:49   Metric: 2
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r2
          State: <Active Int>
          Local AS:    69
          Age: 1:25:49   Metric: 1
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

private__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via lo0.0, selected

```

```
State: <Active Int>
Local AS: 69
Age: 1:31:44
Task: IF
AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0 (1 entry, 1 announced)
  *MPLS Preference: 0
    Next hop type: Receive
    Next-hop reference count: 6
    State: <Active Int>
    Local AS: 69
    Age: 1:31:45 Metric: 1
    Task: MPLS
    Announcement bits (1): 0-KRT
    AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
  *RSVP Preference: 7
    Next hop type: Flood
    Next-hop reference count: 130
    Flood nexthop branches exceed maximum
    Address: 0x8ea65d0

...

800010 (1 entry, 1 announced)
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: via vt-3/2/0.32769, selected
    Label operation: Pop
    State: <Active Int>
    Age: 1:29:30
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
    Label-switched-path green-r1-r3
    Label operation: Push 800012, Push 100096(top)
    Protocol next hop: 10.255.70.103
    Push 800012
    Indirect next hop: 87272e4 1048574
    State: <Active Int>
    Age: 1:29:30 Metric2: 2
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
```

```

abcd::10:255:71:52/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::16/128 (1 entry, 1 announced)
  *MLD Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:43
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.16385, selected
    State: <Active NoReadvrt Int>
    Age: 1:31:44

```

```
Task: IF
AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 1:25:49 Metric2: 1
    AIGP 210
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-1
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-101
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
    Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
```

```

Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
State: <Active Int>
Local AS: 99
Age: 10:21
Task: 12 circuit
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512

```

show route detail (with BGP Multipath)

```
user@host> show route detail
```

```

10.1.1.8/30 (2 entries, 1 announced)
  *BGP Preference: 170/-101
    Next hop type: Router, Next hop index: 262142
    Address: 0x901a010
    Next-hop reference count: 2
    Source: 10.1.1.2
    Next hop: 10.1.1.2 via ge-0/3/0.1, selected
    Next hop: 10.1.1.6 via ge-0/3/0.5
    State: <Active Ext>
    Local AS: 1 Peer AS: 2
    Age: 5:04:43
    Task: BGP_2.10.1.1.2+59955
    Announcement bits (1): 0-KRT
    AS path: 2 I
    Accepted Multipath
    Localpref: 100
    Router ID: 1.1.1.2
  BGP Preference: 170/-101
    Next hop type: Router, Next hop index: 678
    Address: 0x8f97520
    Next-hop reference count: 9
    Source: 10.1.1.6
    Next hop: 10.1.1.6 via ge-0/3/0.5, selected
    State: <NotBest Ext>
    Inactive reason: Not Best in its group - Active preferred
    Local AS: 1 Peer AS: 2
    Age: 5:04:43
    Task: BGP_2.10.1.1.6+58198
    AS path: 2 I
    Accepted MultipathContrib
    Localpref: 100
    Router ID: 1.1.1.3

```

show route exact

Syntax	<code>show route exact <i>destination-prefix</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route exact <i>destination-prefix</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display only the routes that exactly match the specified address or range of addresses.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . <i>destination-prefix</i> —Address or range of addresses. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route exact on page 634 show route exact detail on page 634 show route exact extensive on page 635 show route exact terse on page 635
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route exact

```
user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                  > to 192.168.71.254 via fxp0.0
```

show route exact detail

```
user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
    *Static Preference: 5
```



```

Next-hop reference count: 29
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2d 3:30:26
Task: RT
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

```

show route exact extensive

```

user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:25:18
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

```

show route exact terse

```

user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 207.17.136.0/24  S   5                >192.168.71.254

```

show route export

Syntax	show route export <brief detail> <instance <instance-name> routing-table-name> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route export <brief detail> <instance <instance-name> routing-table-name>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.
Options	<p>none—(Same as brief.) Display standard information about policy-based export for all instances and routing tables on all systems.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <instance-name>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>routing-table-name—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route export inet command).</p>
Required Privilege Level	view
List of Sample Output	show route export on page 637 show route export detail on page 637 show route export instance detail on page 637
Output Fields	Table 65 on page 636 lists the output fields for the show route export command. Output fields are listed in the approximate order in which they appear.

Table 65: show route export Output Fields

Field Name	Field Description	Level of Output
Table or table-name	Name of the routing tables that either import or export routes.	All levels
Routes	Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.	brief none
Export	Whether the table is currently exporting routes to other tables: Y or N (Yes or No).	brief none

Table 65: show route export Output Fields (*continued*)

Field Name	Field Description	Level of Output
Import	Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)	detail
Flags	(instance keyword only) Flags for this feature on this instance: <ul style="list-style-type: none"> config auto-policy—The policy was deduced from the configured IGP export policies. cleanup—Configuration information for this instance is no longer valid. config—The instance was explicitly configured. 	detail
Options	(instance keyword only) Configured option displays the type of routing tables the feature handles: <ul style="list-style-type: none"> unicast—Indicates <i>instance.inet.0</i>. multicast—Indicates <i>instance.inet.2</i>. unicast multicast—Indicates <i>instance.inet.0</i> and <i>instance.inet.2</i>. 	detail
Import policy	(instance keyword only) Policy that route export uses to construct the import-export matrix. Not displayed if the instance type is vrf .	detail
Instance	(instance keyword only) Name of the routing instance.	detail
Type	(instance keyword only) Type of routing instance: forwarding , non-forwarding , or vrf .	detail

Sample Output

show route export

```

user@host> show route export
Table      Export      Routes
inet.0     N            0
black.inet.0 Y           3
red.inet.0 Y            4

```

show route export detail

```

user@host> show route export detail
inet.0                                Routes:    0
black.inet.0                          Routes:    3
  Import: [ inet.0 ]
red.inet.0                             Routes:    4
  Import: [ inet.0 ]

```

show route export instance detail

```

user@host> show route export instance detail
Instance: master                      Type: forwarding
Flags: <config auto-policy> Options: <unicast multicast>
Import policy: [ (ospf-master-from-red || isis-master-from-black) ]

```

Instance: black
Instance: red

Type: non-forwarding
Type: non-forwarding

show route extensive

Syntax	show route extensive <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route extensive <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display extensive information about the active entries in the routing tables.
Options	<p>none—Display all active entries in the routing table.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route extensive on page 644 show route extensive (Access Route) on page 650 show route extensive (Route Reflector) on page 651 show route extensive (FRR and LFA) on page 651 show route extensive (FRR and LFA) on page 652
Output Fields	Table 66 on page 639 describes the output fields for the show route extensive command. Output fields are listed in the approximate order in which they appear.

Table 66: show route extensive Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active). holddown (routes that are in the pending state before being declared inactive). hidden (routes that are not used because of a routing policy).

Table 66: show route extensive Output Fields (*continued*)

Field Name	Field Description
<i>route-destination</i> (entry, announced)	<p>Route destination (for example: 10.0.0.1/24). The entry value is the number of route for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
TSI	Protocol header information.
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[<i>protocol, preference</i>]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
Level	<p>(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>
Route Distinguisher	IP subnet augmented with a 64-bit prefix.

Table 66: show route extensive Output Fields (*continued*)

Field Name	Field Description
Next-hop type	Type of next hop. For a description of possible values for this field, see the Output Field table in the show route detail command.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path <i>lsp-path-name</i>	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.
<i>label-operation</i>	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Indirect next hops	When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.
State	State of the route (a route can be in more than one state). See the Output Field table in the show route detail command.

Table 66: show route extensive Output Fields (*continued*)

Field Name	Field Description
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Inactive reason	<p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> • Active preferred—Currently active route was selected over this route. • Always compare MED—Path with a lower multiple exit discriminator (MED) is available. • AS path—Shorter AS path is available. • Cisco Non-deterministic MED selection—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • Cluster list length—Path with a shorter cluster list length is available. • Forwarding use only—Path is only available for forwarding purposes. • IGP metric—Path through the next hop with a lower IGP metric is available. • IGP metric type—Path with a lower OSPF link-state advertisement type is available. • Interior > Exterior > Exterior via Interior—Direct, static, IGP, or EBGP path is available. • Local preference—Path with a higher local preference value is available. • Next hop address—Path with a lower metric next hop is available. • No difference—Path from a neighbor with a lower IP address is available. • Not Best in its group—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • Number of gateways—Path with a higher number of next hops is available. • Origin—Path with a lower origin code is available. • OSPF version—Path does not support the indicated OSPF version. • RIB preference—Route from a higher-numbered routing table is available. • Route distinguisher—64-bit prefix added to IP subnets to make them unique. • Route metric or MED comparison—Route with a lower metric or MED is available. • Route preference—Route with a lower preference value is available. • Router ID—Path through a neighbor with a lower ID is available. • Unusable path—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • Update source—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.

Table 66: show route extensive Output Fields (*continued*)

Field Name	Field Description
TTL-Action	<p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see show route table.</p>
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • Recorded—The AS path is recorded by the sample process (sampled). • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.

Table 66: show route extensive Output Fields (*continued*)

Field Name	Field Description
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: none or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:34:06
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2

```

```

Next hop: via so-0/3/0.0, selected
State: <Active Int>
Local AS: 69
Age: 1:32:40
Task: IF
Announcement bits (1): 3-Resolve tree 2
AS path: I
OSPF Preference: 10
Next-hop reference count: 1
Next hop: via so-0/3/0.0, selected
State: <Int>
Inactive reason: Route Preference
Local AS: 69
Age: 1:32:40 Metric: 1
Area: 0.0.0.0
Task: OSPF
AS path: I

10.31.1.1/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 7
Interface: so-0/3/0.0
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:32:43
Task: IF
Announcement bits (1): 3-Resolve tree 2
AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kerne 10.31.2.0/30 -> {10.31.1.6}
*OSPF Preference: 10
Next-hop reference count: 9
Next hop: via so-0/3/0.0
Next hop: 10.31.1.6 via ge-3/1/0.0, selected
State: <Active Int>
Local AS: 69
Age: 1:32:19 Metric: 2
Area: 0.0.0.0
Task: OSPF
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kerne 224.0.0.2/32 -> {}
*PIM Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:34:08
Task: PIM Recv
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

```

```
...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.22/32 -> {}
    *IGMP Preference: 0
        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:06
        Task: IGMP
        Announcement bits (2): 0-KRT 3-Resolve tree 2
        AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
State: <FlashAll>
    *RSVP Preference: 7
        Next-hop reference count: 6
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
        Label-switched-path green-r1-r3
        Label operation: Push 100096
        State: <Active Int>
        Local AS: 69
        Age: 1:28:12 Metric: 2
        Task: RSVP
        Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
        AS path: I

10.255.71.238/32 (1 entry, 1 announced)
State: <FlashAll>
    *RSVP Preference: 7
        Next-hop reference count: 6
        Next hop: via so-0/3/0.0 weight 0x1, selected
        Label-switched-path green-r1-r2
        State: <Active Int>
        Local AS: 69
        Age: 1:28:12 Metric: 1
        Task: RSVP
        Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
        AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Active Int>
        Local AS: 69
        Age: 1:34:07
        Task: IF
        AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
```

```

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
    *MPLS Preference: 0
        Next hop type: Receive
        Next-hop reference count: 6
        State: <Active Int>
        Local AS: 69
        Age: 1:34:08 Metric: 1
        Task: MPLS
        Announcement bits (1): 0-KRT
        AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
    *RSVP Preference: 7
        Next hop type: Flood
        Next-hop reference count: 130
        Flood nexthop branches exceed maximum
        Address: 0x8ea65d0

...

800010 (1 entry, 1 announced)
TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: via vt-3/2/0.32769, selected
        Label operation: Pop
        State: <Active Int>
        Age: 1:31:53
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0 /16 -> {indirect(1048574)}
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
        Label-switched-path green-r1-r3
        Label operation: Push 800012, Push 100096(top)
        Protocol next hop: 10.255.70.103
        Push 800012
        Indirect next hop: 87272e4 1048574
        State: <Active Int>
        Age: 1:31:53 Metric2: 2
        Task: Common L2 VC
        Announcement bits (2): 0-KRT 1-Common L2 VC
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0
        Indirect next hops: 1
            Protocol next hop: 10.255.70.103 Metric: 2

```

```

        Push 800012
        Indirect next hop: 87272e4 1048574
        Indirect path forwarding next hops: 1
            Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
        10.255.70.103/32 Originating RIB: inet.3
            Metric: 2                      Node path count: 1
            Forwarding nexthops: 1
                Nexthop: 10.31.1.6 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Active Int>
        Local AS: 69
        Age: 1:34:07
        Task: IF
        AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:07
        Task: IF
        AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
    *PIM Preference: 0
        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:08
        Task: PIM Recv6
        Announcement bits (1): 0-KRT
        AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
    *PIM Preference: 0
        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:08
        Task: PIM Recv6
        Announcement bits (1): 0-KRT
        AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
    *MLD Preference: 0
```

```

Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:34:06
Task: MLD
Announcement bits (1): 0-KRT
AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.16385, selected
    State: <Active NoReadvrt Int>
    Age: 1:34:07
    Task: IF
    AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 1:28:12 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
  *L2VPN Preference: 170/-1
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:34:03 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528

```

```
*L2VPN Preference: 170/-101
Next-hop reference count: 5
Protocol next hop: 10.255.71.52
Indirect next hop: 0 -
State: <Active Int Ext>
Age: 1:34:03 Metric2: 1
Task: green-l2vpn
Announcement bits (1): 1-BGP.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
*L2CKT Preference: 7
Next hop: via so-1/1/2.0 weight 1, selected
Label-switched-path my-lsp
Label operation: Push 100000[0]
Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
State: <Active Int>
Local AS: 99
Age: 10:21
Task: l2 circuit
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512

55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
55.0.0.0/24 -> {Push 300112}
*BGP Preference: 170/-101
Next hop type: Router
Address: 0x925c208
Next-hop reference count: 2
Source: 10.0.0.9
Next hop: 10.0.0.9 via ge-1/2/0.15, selected
Label operation: Push 300112
Label TTL action: prop-ttl
State: <Active Ext>
Local AS: 7019 Peer AS: 13979
Age: 1w0d 23:06:56
AIGP: 25
Task: BGP_13979.10.0.0.9+56732
Announcement bits (1): 0-KRT
AS path: 13979 7018 I
Accepted
Route Label: 300112
Localpref: 100
Router ID: 10.9.9.1
```

show route extensive (Access Route)

```
user@host> show route 13.160.0.102 extensive
```



```

inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
13.160.0.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}
OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern
    *Access Preference: 13
        Next-hop reference count: 78472
        Next hop: 13.160.0.2 via fe-0/0/0.0, selected
        State: <Active Int>
    Age: 12
        Task: RPD Unix Domain Server./var/run/rpd_serv.local
        Announcement bits (2): 0-KRT 1-OSPFv2
        AS path: I

```

show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
    *BGP    Preference: 170/-101
        Source: 192.168.4.214
        Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
        State: <Active Int Ext>
        Local AS: 10458 Peer AS: 10458
        Age: 3:09      Metric: 0      Metric2: 0
        Task: BGP_10458.192.168.4.214+1033
        Announcement bits (2): 0-KRT 4-Resolve inet.0
        AS path: 3944 7777 I <Originator>
        Cluster list: 1.1.1.1
        Originator ID: 10.255.245.88
        Communities: 7777:7777
        Localpref: 100
        Router ID: 4.4.4.4
        Indirect next hops: 1
            Protocol next hop: 207.17.136.192 Metric: 0
            Indirect next hop: 84ac908 40
            Indirect path forwarding next hops: 0
            Next hop type: Discard

```

show route extensive (FRR and LFA)

```

user@host> show route 20.31.2.0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
20.31.2.0/24 (2 entries, 1 announced)
    State: FlashAll

TSI:
KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
    *RSVP    Preference: 7/1
        Next hop type: Router, Next hop index: 1048574
        Address: 0xbbbc010
        Next-hop reference count: 5
        Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299776
        Label TTL action: prop-ttl
        Session Id: 0x201
        Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299792

```

```

Label TTL action: prop-ttl
Session Id: 0x202
State: Active Int
Local AS: 100
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I

```

show route extensive (FRR and LFA)

```

user@host> show route 20:31:2:0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
20.31.2.0/24 (2 entries, 1 announced)
State: FlashAll
TSI:
KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
*RSVP Preference: 7/1
Next hop type: Router, Next hop index: 1048574
Address: 0xbbbc010
Next-hop reference count: 5
Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
Label-switched-path europa-d-to-europa-e
Label operation: Push 299776
Label TTL action: prop-ttl
Session Id: 0x201
Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
Label-switched-path europa-d-to-europa-e
Label operation: Push 299792
Label TTL action: prop-ttl
Session Id: 0x202
State: Active Int
Local AS: 100
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0

```

Task: OSPF
AS path: I

show route flow validation

Syntax	show route flow validation <brief detail> <ip-prefix> <table table-name> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route flow validation <brief detail> <ip-prefix> <table table-name>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display flow route information.
Options	<p>none—Display flow route information.</p> <p>brief detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>ip-prefix—(Optional) IP address for the flow route.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table table-name—(Optional) Display flow route information for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route flow validation inet command).</p>
Required Privilege Level	view
List of Sample Output	show route flow validation on page 655
Output Fields	Table 67 on page 654 lists the output fields for the show route flow validation command. Output fields are listed in the approximate order in which they appear.

Table 67: show route flow validation Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).	All levels
<i>prefix</i>	Route address.	All levels
Active unicast route	Active route in the routing table.	All levels
Dependent flow destinations	Number of flows for which there are routes in the routing table.	All levels

Table 67: show route flow validation Output Fields (*continued*)

Field Name	Field Description	Level of Output
Origin	Source of the route flow.	All levels
Neighbor AS	Autonomous system identifier of the neighbor.	All levels
Flow destination	Number of entries and number of destinations that match the route flow.	All levels
Unicast best match	Destination that is the best match for the route flow.	All levels
Flags	Information about the route flow.	All levels

Sample Output

show route flow validation

```
user@host> show route flow validation
inet.0:
10.0.5.0/24Active unicast route
Dependent flow destinations: 1
Origin: 192.168.224.218, Neighbor AS: 65001
Flow destination (3 entries, 1 match origin)
Unicast best match: 10.0.5.0/24
Flags: SubtreeApex Consistent
```

show route inactive-path

Syntax	<code>show route inactive-path</code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route inactive-path</code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.
Options	none —Display all inactive routes. brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route inactive-path on page 656 show route inactive-path detail on page 657 show route inactive-path extensive on page 658 show route inactive-path terse on page 658
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30      [OSPF/10] 03:57:28, metric 1
> via so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/8          [Direct/0] 04:39:56
> via fxp1.0
```

```

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30      [BGP/170] 04:38:17, localpref 100
                  AS path: 100 I
                  > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route inactive-path detail

```

user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
  OSPF   Preference: 10
         Next-hop reference count: 1
         Next hop: via so-0/3/0.0, selected
         State: <Int>
         Inactive reason: Route Preference
         Local AS:      1
         Age: 3:58:24   Metric: 1
         Area: 0.0.0.0
         Task: OSPF
         AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
  Direct Preference: 0
         Next hop type: Interface
         Next-hop reference count: 1
         Next hop: via fxp1.0, selected
         State: <NotBest Int>
         Inactive reason: No difference
         Age: 4:40:52
         Task: IF
         AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
  BGP   Preference: 170/-101
        Next-hop reference count: 6
        Source: 10.12.80.1

```

```
Next hop: 10.12.80.1 via ge-6/3/2.0, selected
State: <Ext>
Inactive reason: Route Preference
Peer AS: 100
Age: 4:39:13
Task: BGP_100.10.12.80.1+179
AS path: 100 I
Localpref: 100
Router ID: 10.0.0.0
```

show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-path detail on page 657](#).

show route inactive-path terse

```
user@host> show route inactive-path terse
```

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
```

```
Restart Complete
```

```
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	10.12.100.12/30	0	10	1		>so-0/3/0.0	

```
private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	10.0.0.0/8	D	0			>fxp1.0	

```
red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
```

```
Restart Complete
```

```
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	10.12.80.0/30	B	170	100		>10.12.80.1	100 I

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
```

```
Restart Complete
```

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
private1__inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


show route inactive-prefix

Syntax	show route inactive-prefix <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route inactive-prefix <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display inactive route destinations in each routing table.
Options	<p>none—Display all inactive route destination.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route inactive-prefix on page 659 show route inactive-prefix detail on page 659 show route inactive-prefix extensive on page 660 show route inactive-prefix terse on page 660
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-prefix

```

user@host> show route inactive-prefix

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32          [Direct/0] 00:04:54
> via lo0.0

```

show route inactive-prefix detail

```

user@host> show route inactive-prefix detail

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
127.0.0.1/32 (1 entry, 0 announced)
    Direct Preference: 0
    Next hop type: Interface

```

```

Next-hop reference count: 1
Next hop: via 100.0, selected
State: <Hidden Martian Int>
Age: 4:51
Task: IF
AS path: I00:04:54
> via 100.0

```

show route inactive-prefix extensive

The output for the **show route inactive-prefix extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-prefix detail on page 659](#).

show route inactive-prefix terse

```
user@host> show route inactive-prefix terse
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
127.0.0.1/32	D 0			>100.0	

show route instance

Syntax	show route instance <brief detail summary> <instance-name> <logical-system (all <i>logical-system-name</i>)> <operational>
Syntax (EX Series Switches and QFX Series)	show route instance <brief detail summary> <instance-name> <operational>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display routing instance information.
Options	<p>none—(Same as brief) Display standard information about all routing instances.</p> <p>brief detail summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. (These options are not available with the operational keyword.)</p> <p>instance-name—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show route instance cust1 command).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>operational—(Optional) Display operational routing instances.</p>
Required Privilege Level	view
List of Sample Output	show route instance on page 662 show route instance detail (Graceful Restart Complete) on page 663 show route instance detail (Graceful Restart Incomplete) on page 664 show route instance detail (VPLS Routing Instance) on page 666 show route instance operational on page 667 show route instance summary on page 667
Output Fields	Table 68 on page 661 lists the output fields for the show route instance command. Output fields are listed in the approximate order in which they appear.

Table 68: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or <i>instance-name</i>	Name of the routing instance.	All levels

Table 68: show route instance Output Fields (*continued*)

Field Name	Field Description	Level of Output
Operational Routing Instances	(operational keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: forwarding , l2vpn , no-forwarding , vpls , virtual-router , or vrf .	All levels
State	State of the routing instance: active or inactive .	brief detail none
Interfaces	Name of interfaces belonging to this routing instance.	brief detail none
Restart State	Status of graceful restart for this instance: Pending or Complete .	detail
Path selection timeout	Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is 300 .	detail
Tables	Tables (and number of routes) associated with this routing instance.	brief detail none
Route-distinguisher	Unique route distinguisher associated with this routing instance.	detail
Vrf-import	VPN routing and forwarding instance import policy name.	detail
Vrf-export	VPN routing and forwarding instance export policy name.	detail
Vrf-import-target	VPN routing and forwarding instance import target community name.	detail
Vrf-export-target	VPN routing and forwarding instance export target community name.	detail
Fast-reroute-priority	Fast reroute priority setting for a VPLS routing instance: high , medium , or low . The default is low .	detail
Restart State	Restart state: <ul style="list-style-type: none"> Pending:protocol-name—List of protocols that have not yet completed graceful restart for this routing table. Complete—All protocols have restarted for this routing table. 	detail
Primary rib	Primary table for this routing instance.	brief none summary
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Sample Output

show route instance

```

user@host> show route instance
Instance      Type
Primary RIB
master        forwarding
inet.0        16/0/1
iso.0         1/0/0

```

```

mpls.0                                0/0/0
inet6.0                               2/0/0
l2circuit.0                           0/0/0
__juniper_private1__ forwarding
__juniper_private1__.inet.0           12/0/0
__juniper_private1__.inet6.0          1/0/0

```

show route instance detail (Graceful Restart Complete)

```

user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding      State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0              : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3              : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0              : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0             : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0        : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0            : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0        : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
  BGP-INET:
    Router ID: 10.69.103.1
    Type: vrf          State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.103
    Route-distinguisher: 10.255.14.176:103
    Vrf-import: [ BGP-INET-import ]
    Vrf-export: [ BGP-INET-export ]
    Tables:
      BGP-INET.inet.0    : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
  BGP-L:
    Router ID: 10.69.104.1
    Type: vrf          State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.104
    Route-distinguisher: 10.255.14.176:104
    Vrf-import: [ BGP-L-import ]
    Vrf-export: [ BGP-L-export ]
    Tables:
      BGP-L.inet.0      : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
      BGP-L.mpls.0      : 3 routes (3 active, 0 holddown, 0 hidden)
      Restart Complete
  L2VPN:
    Router ID: 0.0.0.0
    Type: l2vpn        State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.512

```

```
Route-distinguisher: 10.255.14.176:512
Vrf-import: [ L2VPN-import ]
Vrf-export: [ L2VPN-export ]
Tables:
  L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
  Restart Complete

LDP:
Router ID: 10.69.105.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.105
Route-distinguisher: 10.255.14.176:105
Vrf-import: [ LDP-import ]
Vrf-export: [ LDP-export ]
Tables:
  LDP.inet.0             : 5 routes (4 active, 0 holddown, 0 hidden)
  Restart Complete

OSPF:
Router ID: 10.69.101.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.101
Route-distinguisher: 10.255.14.176:101
Vrf-import: [ OSPF-import ]
Vrf-export: [ OSPF-export ]
Vrf-import-target: [ target:11111
Tables:
  OSPF.inet.0            : 8 routes (7 active, 0 holddown, 0 hidden)
  Restart Complete

RIP:
Router ID: 10.69.102.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.102
Route-distinguisher: 10.255.14.176:102
Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
  RIP.inet.0             : 6 routes (6 active, 0 holddown, 0 hidden)
  Restart Complete

STATIC:
Router ID: 10.69.100.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.100
Route-distinguisher: 10.255.14.176:100
Vrf-import: [ STATIC-import ]
Vrf-export: [ STATIC-export ]
Tables:
  STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
  Restart Complete
```

show route instance detail (Graceful Restart Incomplete)

```
user@host> show route instance detail
master:
Router ID: 10.255.14.176
```

```

Type: forwarding          State: Active
Restart State: Pending    Path selection timeout: 300
Tables:
  inet.0                  : 17 routes (15 active, 1 holddown, 1 hidden)
Restart Pending: OSPF LDP
  inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: OSPF LDP
  iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
  mpls.0                  : 23 routes (23 active, 0 holddown, 0 hidden)
Restart Pending: LDP VPN
  bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN
  inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
  bgp.l2vpn.0             : 1 routes (1 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf                State: Active
Restart State: Pending    Path selection timeout: 300
Interfaces:
  t3-0/0/0.103
Route-distinguisher: 10.255.14.176:103
Vrf-import: [ BGP-INET-import ]
Vrf-export: [ BGP-INET-export ]
Tables:
  BGP-INET.inet.0         : 6 routes (5 active, 0 holddown, 0 hidden)
Restart Pending: VPN
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf                State: Active
Restart State: Pending    Path selection timeout: 300
Interfaces:
  t3-0/0/0.104
Route-distinguisher: 10.255.14.176:104
Vrf-import: [ BGP-L-import ]
Vrf-export: [ BGP-L-export ]
Tables:
  BGP-L.inet.0            : 6 routes (5 active, 0 holddown, 0 hidden)
Restart Pending: VPN
  BGP-L.mpls.0            : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn              State: Active
Restart State: Pending    Path selection timeout: 300
Interfaces:
  t3-0/0/0.512
Route-distinguisher: 10.255.14.176:512
Vrf-import: [ L2VPN-import ]
Vrf-export: [ L2VPN-export ]
Tables:
  L2VPN.l2vpn.0           : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: VPN L2VPN
LDP:
  Router ID: 10.69.105.1
  Type: vrf                State: Active
Restart State: Pending    Path selection timeout: 300
Interfaces:
  t3-0/0/0.105

```

```
Route-distinguisher: 10.255.14.176:105
Vrf-import: [ LDP-import ]
Vrf-export: [ LDP-export ]
Tables:
  LDP.inet.0          : 5 routes (4 active, 1 holddown, 0 hidden)
Restart Pending: OSPF LDP VPN
OSPF:
Router ID: 10.69.101.1
Type: vrf              State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
  t3-0/0/0.101
Route-distinguisher: 10.255.14.176:101
Vrf-import: [ OSPF-import ]
Vrf-export: [ OSPF-export ]
Tables:
  OSPF.inet.0         : 8 routes (7 active, 1 holddown, 0 hidden)
Restart Pending: OSPF VPN
RIP:
Router ID: 10.69.102.1
Type: vrf              State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
  t3-0/0/0.102
Route-distinguisher: 10.255.14.176:102
Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
  RIP.inet.0          : 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN
STATIC:
Router ID: 10.69.100.1
Type: vrf              State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
  t3-0/0/0.100
Route-distinguisher: 10.255.14.176:100
Vrf-import: [ STATIC-import ]
Vrf-export: [ STATIC-export ]
Tables:
  STATIC.inet.0       : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN
```

show route instance detail (VPLS Routing Instance)

```
user@host> show route instance detail test-vpls
test-vpls:
Router ID: 0.0.0.0
Type: vpls              State: Active
Interfaces:
  lsi.1048833
  lsi.1048832
  fe-0/1/0.513
Route-distinguisher: 10.255.37.65:1
Vrf-import: [ __vrf-import-test-vpls-internal__ ]
Vrf-export: [ __vrf-export-test-vpls-internal__ ]
Vrf-import-target: [ target:300:1 ]
Vrf-export-target: [ target:300:1 ]
Fast-reroute-priority: high
```


Tables:
 test-vpls.l2vpn.0 : 3 routes (3 active, 0 holddown, 0 hidden)

show route instance operational

```
user@host> show route instance operational
Operational Routing Instances:
```

```
master
default
```

show route instance summary

```
user@host> show route instance summary
```

Instance	Type	Primary rib	Active/holddown/hidden
master	forwarding		
		inet.0	15/0/1
		iso.0	1/0/0
		mpls.0	35/0/0
		l3vpn.0	0/0/0
		inet6.0	2/0/0
		l2vpn.0	0/0/0
		l2circuit.0	0/0/0
BGP-INET	vrf		
		BGP-INET.inet.0	5/0/0
		BGP-INET.iso.0	0/0/0
		BGP-INET.inet6.0	0/0/0
BGP-L	vrf		
		BGP-L.inet.0	5/0/0
		BGP-L.iso.0	0/0/0
		BGP-L.mpls.0	4/0/0
		BGP-L.inet6.0	0/0/0
L2VPN	l2vpn		
		L2VPN.inet.0	0/0/0
		L2VPN.iso.0	0/0/0
		L2VPN.inet6.0	0/0/0
		L2VPN.l2vpn.0	2/0/0
LDP	vrf		
		LDP.inet.0	4/0/0
		LDP.iso.0	0/0/0
		LDP.mpls.0	0/0/0
		LDP.inet6.0	0/0/0
		LDP.l2circuit.0	0/0/0
OSPF	vrf		
		OSPF.inet.0	7/0/0
		OSPF.iso.0	0/0/0
		OSPF.inet6.0	0/0/0
RIP	vrf		
		RIP.inet.0	6/0/0
		RIP.iso.0	0/0/0
		RIP.inet6.0	0/0/0
STATIC	vrf		
		STATIC.inet.0	4/0/0
		STATIC.iso.0	0/0/0
		STATIC.inet6.0	0/0/0

show route label

Syntax	<code>show route label <i>label</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route label <i>label</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the routes based on a specified Multiprotocol Label Switching (MPLS) label value.
Options	<i>label</i> —Value of the MPLS label. brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route label on page 668 show route label detail on page 668 show route label extensive on page 669 show route label terse on page 669
Output Fields	For information about output fields, see the output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route label

```
user@host> show route label 100016

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
100016          *[VPN/170] 03:25:41
                  > to 10.12.80.1 via ge-6/3/2.0, Pop
```

show route label detail

```
user@host> show route label 100016 detail

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
100016 (1 entry, 1 announced)
```

```

*VPN      Preference: 170
          Next-hop reference count: 2
          Source: 10.12.80.1
          Next hop: 10.12.80.1 via ge-6/3/2.0, selected
          Label operation: Pop
          State: <Active Int Ext>
          Local AS:      1
          Age: 3:23:31
          Task: BGP.0.0.0.0+179
          Announcement bits (1): 0-KRT
          AS path: 100 I
          Ref Cnt: 2

```

show route label extensive

The output for the `show route label extensive` command is identical to that of the `show route label detail` command. For sample output, see [show route label detail on page 668](#).

show route label terse

```

user@host> show route label 100016 terse

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 100016          V 170                >10.12.80.1

```

show route label-switched-path

Syntax	<code>show route label-switched-path <i>path-name</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route label-switched-path <i>path-name</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the routes used in an MPLS label-switched path (LSP).
Options	brief detail extensive terse —(Optional) Display the specified level of output. <i>path-name</i> —LSP tunnel name. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route label-switched-path on page 670
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route label-switched-path

```
user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          [MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32          *[MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32          *[MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
4.4.4.4/32          *[MPLS/7] 00:00:06, metric 0
                    to 111.222.1.9 via s0-0/0/0, label-switched-path abc
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
                    to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
```

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

show route martians

Syntax	show route martians <logical-system (all <i>logical-system-name</i>)> <table <i>routing-table-name</i> >
Syntax (EX Series Switches)	show route martians <table <i>routing-table-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the martian (invalid and ignored) entries associated with each routing table.
Options	<p>none—Display standard information about route martians for all routing tables.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display information about route martians for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route martians table inet command).</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Martian Addresses
List of Sample Output	show route martians on page 672
Output Fields	Table 69 on page 672 lists the output fields for the show route martians command. Output fields are listed in the approximate order in which they appear

Table 69: show route martians Output Fields

Field Name	Field Description
<i>table-name</i>	Name of the route table in which the route martians reside.
<i>destination-prefix</i>	Route destination.
<i>match value</i>	Route match parameter.
<i>status</i>	Status of the route: allowed or disallowed .

Sample Output

show route martians

```
user@host> show route martians
```

```
inet.0:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.1:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed

inet.2:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.3:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

...

inet6.0:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.1:
    ::1/128 exact -- disallowed

inet6.2:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.3:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

...
```

show route next-hop

Syntax	<code>show route next-hop <i>next-hop</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route next-hop <i>next-hop</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that are being sent to the specified next-hop address.
Options	brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>next-hop</i> —Next-hop address.
Required Privilege Level	view
List of Sample Output	show route next-hop on page 674 show route next-hop detail on page 675 show route next-hop extensive on page 677 show route next-hop terse on page 678
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route next-hop

```
user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
172.16.0.0/12    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
192.168.0.0/16   *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
192.168.102.0/23 *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
```



```

207.17.136.0/24    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
207.17.136.192/32 *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route next-hop detail

```

user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 36
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 1
        Age: 6:27:41
        Task: RT
        Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
        AS path: I

10.209.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 36
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 1
        Age: 6:27:41
        Task: RT
        Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
        AS path: I

172.16.0.0/12 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 36
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 1
        Age: 6:27:41
        Task: RT
        Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
        AS path: I

192.168.0.0/16 (1 entry, 1 announced)

```

```
*Static Preference: 5
  Next-hop reference count: 36
  Next hop: 192.168.71.254 via fxp0.0, selected
  State: <Active NoReadvrt Int Ext>
  Local AS: 1
  Age: 6:27:41
  Task: RT
  Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
  AS path: I

192.168.102.0/23 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.192/32 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop extensive

```

user@host> show route next-hop 192.168.71.254 extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 22
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 69
        Age: 2:02:28
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 22
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 69
        Age: 2:02:28
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 22
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 69
        Age: 2:02:28
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

192.168.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 22
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 69
        Age: 2:02:28
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

192.168.102.0/23 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
    *Static Preference: 5

```

```

Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
*Static Preference: 5
Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
*Static Preference: 5
Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route next-hop terse

```

user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.10.0.0/16      S  5          5          >192.168.71.254
* 10.209.0.0/16     S  5          5          >192.168.71.254
* 172.16.0.0/12     S  5          5          >192.168.71.254

```

```
* 192.168.0.0/16      S   5                >192.168.71.254
* 192.168.102.0/23   S   5                >192.168.71.254
* 207.17.136.0/24    S   5                >192.168.71.254
* 207.17.136.192/32  S   5                >192.168.71.254

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route no-community

Syntax	show route no-community <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route no-community <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are not associated with any community.
Options	<p>none—(Same as brief) Display the route entries in each routing table that are not associated with any community.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route no-community on page 680 show route no-community detail on page 681 show route no-community extensive on page 681 show route no-community terse on page 682
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route no-community

```

user@host> show route no-community
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:36:27
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:36:27
                  > to 192.168.71.254 via fxp0.0
10.255.71.52/32  *[Direct/0] 00:36:27
                  > via lo0.0
10.255.71.63/32  *[OSPF/10] 00:04:39, metric 1
                  > to 35.1.1.2 via ge-3/1/0.0
10.255.71.64/32  *[OSPF/10] 00:00:08, metric 2
                  > to 35.1.1.2 via ge-3/1/0.0
10.255.71.240/32 *[OSPF/10] 00:05:04, metric 2
                  via so-0/1/2.0

```

```

> via so-0/3/2.0
10.255.71.241/32 * [OSPF/10] 00:05:14, metric 1
> via so-0/1/2.0
10.255.71.242/32 * [OSPF/10] 00:05:19, metric 1
> via so-0/3/2.0
12.1.1.0/24 * [OSPF/10] 00:05:14, metric 2
> via so-0/3/2.0
14.1.1.0/24 * [OSPF/10] 00:00:08, metric 3
> to 35.1.1.2 via ge-3/1/0.0
via so-0/1/2.0
via so-0/3/2.0
16.1.1.0/24 * [OSPF/10] 00:05:14, metric 2
> via so-0/1/2.0
.....

```

show route no-community detail

```

user@host> show route no-community detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

....

```

show route no-community extensive

```

user@host> show route no-community extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:

```

```
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I
```

show route no-community terse

```
user@host> show route no-community terse
```

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.10.0.0/16	S	5			>192.168.71.254	
*	10.209.0.0/16	S	5			>192.168.71.254	
*	10.255.71.52/32	D	0			>lo0.0	
*	10.255.71.63/32	0	10	1		>35.1.1.2	
*	10.255.71.64/32	0	10	2		>35.1.1.2	
*	10.255.71.240/32	0	10	2		so-0/1/2.0	
						>so-0/3/2.0	
*	10.255.71.241/32	0	10	1		>so-0/1/2.0	
*	10.255.71.242/32	0	10	1		>so-0/3/2.0	
*	12.1.1.0/24	0	10	2		>so-0/3/2.0	
*	14.1.1.0/24	0	10	3		>35.1.1.2	
						so-0/1/2.0	
						so-0/3/2.0	
*	16.1.1.0/24	0	10	2		>so-0/1/2.0	
...							

show route protocol

Syntax	<pre>show route protocol <i>protocol</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show route protocol <i>protocol</i> <brief detail extensive terse></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>ospf2 and ospf3 options introduced in Junos OS Release 9.2.</p> <p>ospf2 and ospf3 options introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>flow option introduced in Junos OS Release 10.0.</p> <p>flow option introduced in Junos OS Release 10.0 for EX Series switches.</p>
Description	Display the route entries in the routing table that were learned from a particular protocol.
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>protocol</i>—Protocol from which the route was learned:</p> <ul style="list-style-type: none"> • access—Access route for use by DHCP application • access-internal—Access-internal route for use by DHCP application • aggregate—Locally generated aggregate route • arp—Route learned through the Address Resolution Protocol • atmvpn—Asynchronous Transfer Mode virtual private network • bgp—Border Gateway Protocol • ccc—Circuit cross-connect • direct—Directly connected route • dvmrp—Distance Vector Multicast Routing Protocol • esis—End System-to-Intermediate System • flow—Locally defined flow-specification route • frr—Precomputed protection route or backup route used when a link goes down • isis—Intermediate System-to-Intermediate System • ldp—Label Distribution Protocol • l2circuit—Layer 2 circuit • l2vpn—Layer 2 virtual private network

- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



NOTE: EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level	view
List of Sample Output	show route protocol access on page 685 show route protocol access-internal extensive on page 685 show route protocol arp on page 685 show route protocol bgp on page 686 show route protocol bgp detail on page 686 show route protocol bgp extensive on page 686 show route protocol bgp terse on page 687 show route protocol direct on page 687 show route protocol frr on page 688 show route protocol l2circuit detail on page 688 show route protocol l2vpn extensive on page 689 show route protocol ldp on page 690 show route protocol ldp extensive on page 690 show route protocol ospf (Layer 3 VPN) on page 691 show route protocol ospf detail on page 692 show route protocol rip on page 692 show route protocol rip detail on page 692 show route protocol ripng table inet6 on page 693 show route protocol static detail on page 693

Output Fields For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
    *Access-internal Preference: 12
        Next-hop reference count: 200000
        Next hop: 13.160.0.2 via fe-0/0/0.0, selected
        State: <Active Int>
    Age: 36
        Task: RPD Unix Domain Server./var/run/rpd_serv.local
        Announcement bits (1): 0-KRT
        AS path: I
```

show route protocol arp

```
user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.4/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.5/32      [ARP/4294967293] 00:04:32, from 20.20.1.1
                  Unusable
20.20.1.6/32      [ARP/4294967293] 00:04:34, from 20.20.1.1
                  Unusable
20.20.1.7/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.8/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.9/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.10/32     [ARP/4294967293] 00:04:35, from 20.20.1.1
```

```

Unusable
20.20.1.11/32      [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
20.20.1.12/32      [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
20.20.1.13/32      [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
...

```

show route protocol bgp

```

user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21      *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
AS path: 10458 14203 2914 4788 4788 I
> to 192.168.167.254 via fxp0.0

```

show route protocol bgp detail

```

user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24      (1 entry, 1 announced)
    *BGP      Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006436
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS: 69 Peer AS: 10458
                Age: 6d 10:42:42 Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1
    AS path: 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400
    Accepted
    Localpref: 100
    Router ID: 207.17.136.192

```

show route protocol bgp extensive

```

user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
  Nexthop: Self
  AS path: [69] 10458 14203 2914 4788 4788 I
  Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1
    *BGP      Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006502
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324

```

```

Next hop: 192.168.167.254 via fxp0.0, selected
Protocol next hop: 192.168.69.71
Indirect next hop: 8e166c0 342
State: <Active Ext>
Local AS: 69 Peer AS: 10458
Age: 6d 10:44:45 Metric2: 0
Task: BGP_10458.192.168.69.71+179
Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1
AS path: 10458 14203 2914 4788 4788 I
Communities: 2914:410 2914:2403 2914:3400
Accepted
Localpref: 100
Router ID: 207.17.136.192
Indirect next hops: 1
  Protocol next hop: 192.168.69.71
  Indirect next hop: 8e166c0 342
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 192.168.167.254 via fxp0.0
  192.168.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
    Nexthop: 192.168.167.254 via fxp0.0

```

show route protocol bgp terse

```

user@host> show route protocol bgp 192.168.64.0/21 terse

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
192.168.64.0/21   B 170      100          >100.1.3.2    10023 21 I

```

show route protocol direct

```

user@host> show route protocol direct

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

8.8.8.0/24          *[Direct/0] 17w0d 10:31:49
> via fe-1/3/1.0
10.255.165.1/32     *[Direct/0] 25w4d 04:13:18
> via lo0.0
30.30.30.0/24       *[Direct/0] 17w0d 23:06:26
> via fe-1/3/2.0
192.168.164.0/22    *[Direct/0] 25w4d 04:13:20
> via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
*[Direct/0] 25w4d 04:13:21
> via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

abcd::10:255:165:1/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0

```

show route protocol frr

```

user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.3 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.4 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32      *[FRR/200] 00:05:35, from 20.20.1.1
                  > to 20.20.1.5 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32      *[FRR/200] 00:05:37, from 20.20.1.1
                  > to 20.20.1.6 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.7 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.8 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.9 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32     *[FRR/200] 00:05:38, from 20.20.1.1
...

```

show route protocol l2circuit detail

```

user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop: via ge-2/0/0.0, selected
        Label operation: Pop      Offset: 4
        State: <Active Int>
        Local AS: 99
        Age: 9:52
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop: via so-1/1/2.0 weight 1, selected
        Label-switched-path my-lsp

```

```

Label operation: Push 100000, Push 100000(top)[0] Offset: -4
Protocol next hop: 10.245.255.63
Push 100000 Offset: -4
  Indirect next hop: 86af0c0 298
State: <Active Int>
Local AS: 99
Age: 9:52
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

show route protocol l2vpn extensive

```

user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
  *L2VPN Preference: 7
    Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
    Label operation: Pop Offset: 4
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
  *L2VPN Preference: 7
    Next hop: via so-0/0/1.0, selected
    Label operation: Push 800000 Offset: -4
    Protocol next hop: 10.255.14.220
    Push 800000 Offset: -4
    Indirect next hop: 85142a0 288
    State: <Active Int>

```

```
Local AS:    69
Age: 7:48
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I
Communities: target:69:1 Layer2-info: encaps:PPP,
control flags:2, mtu: 0
```

show route protocol ldp

```
user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100080            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Swap 100000
```

show route protocol ldp extensive

```
user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          Label operation: Push 100000
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

192.168.17.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I
```



```
private1__inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
```

```
100064 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
```

```
*LDP      Preference: 9
           Next-hop reference count: 2
           Next hop: via t1-4/0/0.0, selected
           State: <Active Int>
           Local AS: 65500
           Age: 1d 23:03:58      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
           Prefixes bound to route: 192.168.17.1/32
```

```
100064(S=0) (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
```

```
*LDP      Preference: 9
           Next-hop reference count: 2
           Next hop: via t1-4/0/0.0, selected
           Label operation: Pop
           State: <Active Int>
           Local AS: 65500
           Age: 1d 23:03:58      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
```

```
100080 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
```

```
*LDP      Preference: 9
           Next-hop reference count: 2
           Next hop: via t1-4/0/0.0, selected
           Label operation: Swap 100000
           State: <Active Int>
           Local AS: 65500
           Age: 1d 23:03:58      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
           Prefixes bound to route: 192.168.16.1/32
```

show route protocol ospf (Layer 3 VPN)

```
user@host> show route protocol ospf
```

```
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
10.39.1.4/30      *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.39.1.8/30      [OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
10.255.14.171/32  *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.255.14.179/32  *[OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
```

```
224.0.0.5/32      *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30     [OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
10.255.14.173/32  *[OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
224.0.0.5/32      *[OSPF/10] 20:26:20, metric 1
```

show route protocol ospf detail

```
user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
    OSPF    Preference: 10
            Nexthop: via so-0/2/2.0, selected
            State: <Int>
            Inactive reason: Route Preference
            Age: 6:25      Metric: 1
            Area: 0.0.0.0
            Task: VPN-AB-OSPF
            AS path: I
            Communities: Route-Type:0.0.0.0:1:0

...
```

show route protocol rip

```
user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32  *[RIP/100] 20:24:34, metric 2
                  > to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32     *[RIP/100] 00:03:59, metric 1
```

show route protocol rip detail

```
user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
    *RIP    Preference: 100
            Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
            State: <Active Int>
            Age: 20:25:02  Metric: 2
            Task: VPN-AB-RIPv2
            Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
            AS path: I
            Route learned from 10.39.1.22 expires in 96 seconds
```

show route protocol ripng table inet6

```

user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0

```

show route protocol static detail

```

user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.13.10.0/23 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified

```

Task: RT
Announcement bits (1): 0-KRT
AS path: I

show route range

Syntax	<pre>show route range <brief detail extensive terse> <destination-prefix> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show route range <brief detail extensive terse> <destination-prefix></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display routing table entries using a prefix range.
Options	<p>none—Display standard information about all routing table entries using a prefix range.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>destination-prefix—(Optional) Destination and prefix mask for the range.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	<p>show route range on page 695</p> <p>show route range destination-prefix on page 696</p> <p>show route range detail on page 696</p> <p>show route range extensive on page 697</p> <p>show route range terse on page 698</p>
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route range

```
user@host> show route range

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.255.71.14/32  *[Direct/0] 00:30:01
                  > via lo0.0
```

```
172.16.0.0/12      *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
192.168.0.0/16     *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:30:01
                  > via fxp0.0
192.168.71.14/32   *[Local/0] 00:30:01
                  Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
...
```

show route range destination-prefix

```
user@host> show route range 192.168.0.0
```

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.168.0.0/16     *[Static/5] 00:31:14
                  > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:31:14
                  > via fxp0.0
192.168.71.14/32   *[Local/0] 00:31:14
                  Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:31:14
                  > to 192.168.71.254 via fxp0.0
```

show route range detail

```
user@host> show route range detail
```

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
```

```
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I
```

```
10.209.0.0/16 (1 entry, 1 announced)
```

```
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I
```

```
10.255.71.14/32 (1 entry, 0 announced)
```

```
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Age: 30:05
    Task: IF
```

```

AS path: I
172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

```

...

show route range extensive

```

user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Age: 30:17
    Task: IF
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17

```

Task: RT
 Announcement bits (1): 0-KRT
 AS path: I

...

show route range terse

user@host> show route range terse

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
 + = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
* 10.10.0.0/16	S 5			>192.168.71.254	
* 10.209.0.0/16	S 5			>192.168.71.254	
* 10.255.71.14/32	D 0			>lo0.0	
* 172.16.0.0/12	S 5			>192.168.71.254	
* 192.168.0.0/16	S 5			>192.168.71.254	
* 192.168.64.0/21	D 0			>fxp0.0	
* 192.168.71.14/32	L 0			Local	
* 192.168.102.0/23	S 5			>192.168.71.254	
* 207.17.136.0/24	S 5			>192.168.71.254	
* 207.17.136.192/32	S 5			>192.168.71.254	

__juniper_private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
* 10.0.0.0/8	D 0			>fxp2.0	
	D 0			>fxp1.0	
* 10.0.0.4/32	L 0			Local	

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152					
*	D 0			>lo0.0	

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
abcd::10:255:71:14/128					
*	D 0			>lo0.0	
fe80::280:42ff:fe11:226f/128					
*	D 0			>lo0.0	

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
fe80::280:42ff:fe11:226f/128					
*	D 0			>lo0.16385	

show route receive-protocol

Syntax	show route receive-protocol <i>protocol neighbor-address</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>	
Syntax (EX Series Switches)	show route receive-protocol <i>protocol neighbor-address</i> <brief detail extensive terse>	
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.	
Description	Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.	
Options	brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>protocol neighbor-address</i> —Protocol transmitting the route (bgp , dvmrp , msdp , pim , rip , or ripng) and address of the neighboring router from which the route entry was received.	
Additional Information	The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.	
Required Privilege Level	view	
List of Sample Output	show route receive-protocol bgp on page 702 show route receive-protocol bgp extensive on page 702 show route receive-protocol bgp table extensive on page 702 show route receive-protocol bgp logical-system extensive on page 703 show route receive-protocol bgp detail (Layer 2 VPN) on page 704 show route receive-protocol bgp extensive (Layer 2 VPN) on page 704 show route receive-protocol bgp (Layer 3 VPN) on page 705 show route receive-protocol bgp detail (Layer 3 VPN) on page 705 show route receive-protocol bgp extensive (Layer 3 VPN) on page 706	
Output Fields	Table 70 on page 699 describes the output fields for the show route receive-protocol command. Output fields are listed in the approximate order in which they appear.	

Table 70: show route receive-protocol Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table—for example, inet.0.	All levels
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.	All levels

Table 70: show route receive-protocol Output Fields (*continued*)

Field Name	Field Description	Level of Output
number routes	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holddown (routes that are in pending state before being declared inactive) • hidden (routes that are not used because of a routing policy) 	All levels
Prefix	Destination prefix.	none brief
MED	Multiple exit discriminator value included in the route.	none brief
destination-prefix (entry, announced)	Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination.	detail extensive
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.	detail extensive
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.	detail extensive
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.	detail extensive
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	All levels
Localpref or Lclpref	Local preference value included in the route.	All levels

Table 70: show route receive-protocol Output Fields (*continued*)

Field Name	Field Description	Level of Output
AS path	<p>Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893. • []—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>	All levels
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	detail extensive
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	detail extensive
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.	detail extensive
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	detail extensive
Attrset AS	Number, local preference, and path of the AS that originated the route. These values are stored in the Attrset attribute at the originating routing device.	detail extensive
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	detail extensive
control flags	Control flags: none or Site Down .	detail extensive
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	detail extensive

Sample Output

show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
10.22.1.0/24     10.255.245.215    0        100      I
10.22.2.0/24     10.255.245.215    0        100      I
```

show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
1.1.1.0/24 (1 entry, 1 announced)
  Next hop: 10.0.50.3
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
165.3.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
165.4.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
195.1.2.0/24 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp table extensive

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
  Nexthop: 207.17.136.29
  Localpref: 100
  AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
  AS path: AS4 PA[2]: 33437 393219
  AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
  Communities: 2914:420
```

show route receive-protocol bgp logical-system extensive

```

user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

10.0.0.8/30 (2 entries, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 44.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300096
  Nexthop: 10.0.0.9
  AS path: 13979 I
  AIGP: 203

* 55.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300112
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I
  AIGP: 25

* 66.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300144
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I

* 99.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300160
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I

```

show route receive-protocol bgp detail (Layer 2 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

show route receive-protocol bgp extensive (Layer 2 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100

```

```

AS path: I
Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

show route receive-protocol bgp (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171
inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED      Lclpref AS path
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.179/32 10.255.14.171          2    100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.177/32 10.255.14.171          100 I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.14.171:300:10.255.14.177/32
                  10.255.14.171          100 I
10.255.14.171:100:10.255.14.179/32
                  10.255.14.171          2    100 I
10.255.14.171:200:10.255.14.175/32
                  10.255.14.171          100 2 I

```

show route receive-protocol bgp detail (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

```

* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

show route receive-protocol bgp extensive (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
  1.1.1.0/24 (1 entry, 1 announced)
    Nexthop: 10.0.50.3
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  165.3.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
  165.4.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  195.1.2.0/24 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)

```


show route resolution

Syntax	<pre>show route resolution <brief detail extensive summary> <index <i>index</i>> <logical-system (all <i>logical-system-name</i>)> <prefix> <table <i>routing-table-name</i>> <unresolved></pre>
Syntax (EX Series Switches)	<pre>show route resolution <brief detail extensive summary> <index <i>index</i>> <prefix> <table <i>routing-table-name</i>> <unresolved></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.</p>
Options	<p>none—Display standard information about all entries in the next-hop resolution database.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>index <i>index</i>—(Optional) Show the index of the resolution tree.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix <i>network/destination-prefix</i>—(Optional) Display database entries for the specified address.</p> <p>table <i>routing-table-name</i>—(Optional) Display information about a particular routing table (for example, inet.0) where policy-based export is currently enabled.</p> <p>unresolved—(Optional) Display routes that could not be resolved.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Route Resolution on PE Routers
List of Sample Output	<p>show route resolution detail on page 708</p> <p>show route resolution summary on page 709</p> <p>show route resolution unresolved on page 709</p>

Output Fields Table 71 on page 708 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

Table 71: show route resolution Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.
Tree index	Tree index identifier.
Nodes	Number of nodes in the tree.
Reference count	Number of references made to the next hop.
Contributing routing tables	Routing tables used for next-hop resolution.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3 , this field indicates which routing table, inet.0 or inet.3 , provided the best path for a particular prefix.
Metric	Metric associated with the forwarding next hop.
Node path count	Number of nodes in the path.
Forwarding next hops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route resolution detail

```

user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0

```

```

Node path count: 1
Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
Metric: 2 Node path count: 1
Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1

```

show route resolution summary

```

user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
Contributing routing tables: :voice.inet.0 :voice.inet.3
Tree Index: 2, Nodes 2, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 3, Nodes 43, Reference Count 1
Contributing routing tables: inet.0 inet.3

```

show route resolution unresolved

```

user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
  Protocol Nexthop: 10.255.71.238 Push 800000
  Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
  Protocol Nexthop: 10.255.70.103 Push 800008
  Indirect nexthop: 0 -
Tree Index 2

```

show route snooping

Syntax	<code>show route snooping</code> <code><brief detail extensive terse></code> <code><all></code> <code><best address/prefix></code> <code><exact address></code> <code><range prefix-range></code> <code><summary></code> <code><table table-name></code>
Release Information	Command introduced in Junos OS Release 8.5. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that were learned from snooping.
Options	<p>none—Display the entries in the routing table that were learned from snooping.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>all—(Optional) Display all entries, including hidden entries.</p> <p>best address/prefix—(Optional) Display the longest match for the provided address and optional prefix.</p> <p>exact address/prefix—(Optional) Display exact matches for the provided address and optional prefix.</p> <p>range prefix-range—(Optional) Display information for the provided address range.</p> <p>summary—(Optional) Display route snooping summary statistics.</p> <p>table table-name—(Optional) Display information for the named table.</p>
Required Privilege Level	view
List of Sample Output	show route snooping detail on page 710
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route snooping detail

```
user@host> show route snooping detail
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
224.0.0.2/32 (1 entry, 1 announced)
  *IGMP    Preference: 0
           Next hop type: MultiRecv
           Next-hop reference count: 4
           State: <Active NoReadvrt Int>
```

```

Age: 2:24
Task: IGMP
Announcement bits (1): 0-KRT
AS path: I

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
    Next hop type: MultiRecv
    Next-hop reference count: 4
    State: <Active NoReadvrt Int>
    Age: 2:24
    Task: IGMP
    Announcement bits (1): 0-KRT
    AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4), Next hop index: 1048584
    Next-hop reference count: 4
    State: <Active Int>
    Age: 2:24
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.2.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.3.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.4.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.5.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113

```

```
State: <Active Int>
Age: 1:58
Task: MC
Announcement bits (1): 0-KRT
AS path: I

225.0.0.6.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.7.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.9.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.10.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.1.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.2.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
```

```
Age: 8
Task: MC
Announcement bits (1): 0-KRT
AS path: I

226.0.0.4.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.8.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.10.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:56
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.1.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.2.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.3.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:16
```

```

Task: MC
Announcement bits (1): 0-KRT
AS path: I

227.0.0.4.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.5.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.7.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.8.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.10.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.1.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC

```



```
Announcement bits (1): 0-KRT
AS path: I

228.0.0.2.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:18
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.7.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:11
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.8.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.9.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 8
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.10.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.3.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
```

```
AS path: I

229.0.0.4.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.5.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 9
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.6.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.7.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.8.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.9.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

```
229.0.0.10.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

show route source-gateway

Syntax	<code>show route source-gateway <i>address</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route source-gateway <i>address</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that were learned from a particular address. The Source field in the <code>show route detail</code> command output lists the source for each route, if known.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . <i>address</i> —IP address of the system. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route source-gateway on page 718 show route source-gateway detail on page 719 show route source-gateway extensive on page 721
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route source-gateway

```
user@host> show route source-gateway 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
```

```

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1__inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

```

show route source-gateway detail

```

user@host> show route source-gateway 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
    *BGP      Preference: 170/-101

```

```
Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:14:00 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-green-l2vpn
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
*BGP Preference: 170/-1
Route Distinguisher: 10.255.70.103:2
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:14:00 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-red-l2vpn
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down, mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:14:00 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS, control
flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
```

```

Secondary Tables: green.l2vpn.0
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
  *BGP Preference: 170/-1
    Route Distinguisher: 10.255.70.103:2
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:14:00 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    AS path: I
    Communities: target:11111:2 Layer2-info: encaps:VPLS,
    control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8
    Localpref: 0
    Router ID: 10.255.70.103
    Secondary Tables: red.l2vpn.0

```

show route source-gateway extensive

```

user@host> show route source-gateway 10.255.70.103 extensive
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:15:24 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103

```

Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)

```
*BGP      Preference: 170/-1
          Route Distinguisher: 10.255.70.103:2
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Secondary Active Int Ext>
          Local AS:    69 Peer AS:    69
          Age: 12:15:24  Metric2: 1
          Task: BGP_69.10.255.70.103+179
          Announcement bits (1): 0-red-l2vpn
          AS path: I
          Communities: target:11111:2 Layer2-info: encaps:VPLS,
          control flags:Site-Down, mtu: 0
          Label-base: 800016, range: 8
          Localpref: 0
          Router ID: 10.255.70.103
          Primary Routing Table bgp.l2vpn.0
```

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)

```
*BGP      Preference: 170/-101
          Route Distinguisher: 10.255.70.103:1
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Active Int Ext>
          Local AS:    69 Peer AS:    69
          Age: 12:15:24  Metric2: 1
          Task: BGP_69.10.255.70.103+179
          AS path: I
          Communities: target:11111:1 Layer2-info: encaps:VPLS,
          control flags:, mtu: 0
          Label-base: 800008, range: 8
          Localpref: 100
          Router ID: 10.255.70.103
          Secondary Tables: green.l2vpn.0
          Indirect next hops: 1
            Protocol next hop: 10.255.70.103 Metric: 2
            Indirect next hop: 2 no-forward
            Indirect path forwarding next hops: 1
          Next hop:      via so-0/3/0.0 weight 0x1
            10.255.70.103/32 Originating RIB: inet.3
              Metric: 2                      Node path count: 1
              Forwarding nexthops: 1
                Nexthop: via so-0/3/0.0
```

10.255.70.103:2:3:1/96 (1 entry, 0 announced)

```
*BGP      Preference: 170/-1
          Route Distinguisher: 10.255.70.103:2
          Next-hop reference count: 7
          Source: 10.255.70.103
```



```
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Secondary Tables: red.12vpn.0
Indirect next hops: 1
    Protocol next hop: 10.255.70.103 Metric: 2
    Indirect next hop: 2 no-forward
    Indirect path forwarding next hops: 1
Next hop:      via so-0/3/0.0 weight 0x1
    10.255.70.103/32 Originating RIB: inet.3
    Metric: 2 Node path count: 1
    Forwarding nexthops: 1
    Nexthop: via so-0/3/0.0
```

show route summary

Syntax	show route summary <logical-system (all <i>logical-system-name</i>)> <table <i>routing-table-name</i> >
Syntax (EX Series Switches)	show route summary
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Display summary statistics about the entries in the routing table.</p> <p>CPU utilization might increase while the device learns routes. We recommend that you use the show route summary command after the device learns and enters the routes into the routing table. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the show route summary command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the command-line interface (CLI).</p>
Options	<p>none—Display summary statistics about the entries in the routing table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display summary statistics for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route summary table inet command). If you only want to display statistics for a specific routing table, make sure to enter the exact name of that routing table.</p>
Required Privilege Level	view
List of Sample Output	show route summary on page 725 show route summary table on page 726 show route summary table (with Route Limits Configured for the Routing Table) on page 726
Output Fields	Table 72 on page 724 lists the output fields for the show route summary command. Output fields are listed in the approximate order in which they appear.

Table 72: show route summary Output Fields

Field Name	Field Description
Router ID	Address of the local routing device.
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).

Table 72: show route summary Output Fields (*continued*)

Field Name	Field Description
destinations	Number of destinations for which there are routes in the routing table.
routes	Number of routes in the routing table: <ul style="list-style-type: none"> • active—Number of routes that are active. • holddown—Number of routes that are in the hold-down state before being declared inactive. • hidden—Number of routes that are not used because of routing policy.
Limit/Threshold	Displays the configured route limits for the routing table set with the maximum-prefixes and the maximum-paths statements. If you do not configure route limits for the routing table, the show output does not display this information. <ul style="list-style-type: none"> • destinations—The first number represents the maximum number of route prefixes installed in the routing table. The second number represents the number of route prefixes that trigger a warning message. • routes—The first number represents the maximum number of routes. The second number represents the number of routes that trigger a warning message.
Direct	Routes on the directly connected network.
Local	Local routes.
protocol-name	Name of the protocol from which the route was learned. For example, OSPF, RSVP, and Static.

Sample Output

show route summary

```

user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
    Direct:      6 routes,      5 active
    Local:       4 routes,      4 active
    OSPF:        5 routes,      4 active
    Static:      7 routes,      7 active
    IGMP:        1 routes,      1 active
    PIM:         2 routes,      2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
    RSVP:        2 routes,      2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
    Direct:      1 routes,      1 active

```

```

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
      MPLS:      3 routes,      3 active
      VPLS:      4 routes,      2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
      Direct:     2 routes,      2 active
      PIM:        2 routes,      2 active
      MLD:        1 routes,      1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
      BGP:        2 routes,      2 active
      L2VPN:      2 routes,      2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
      BGP:        2 routes,      2 active
      L2VPN:      1 routes,      1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
      BGP:        4 routes,      4 active

```

show route summary table

```

user@host> show route summary table inet
Router ID: 192.168.0.1

inet.0: 32 destinations, 34 routes (31 active, 0 holddown, 1 hidden)
      Direct:     6 routes,      5 active
      Local:      9 routes,      9 active
      OSPF:       3 routes,      1 active
      Static:    13 routes,     13 active
      IGMP:       1 routes,      1 active
      PIM:        2 routes,      2 active

inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Multicast:   1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
      Local:      1 routes,      1 active
      PIM:        2 routes,      2 active

inet6.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Multicast:   1 routes,      1 active

```

show route summary table (with Route Limits Configured for the Routing Table)

```

user@host> show route summary table VPN-A.inet.0
Autonomous system number: 100
Router ID: 10.255.182.142

VPN-A.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
Limit/Threshold: 2000/200 destinations 20/12 routes
      Direct:     2 routes,      2 active
      Local:      1 routes,      1 active
      OSPF:       4 routes,      3 active
      BGP:        4 routes,      4 active

```

IGMP:	1 routes,	1 active
PIM:	2 routes,	2 active

show route table

Syntax	<code>show route table <i>routing-table-name</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route table <i>routing-table-name</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in a particular routing table.
Options	brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>routing-table-name</i> —Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route table inet command).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show route summary on page 724
List of Sample Output	show route table bgp.l2.vpn on page 729 show route table bgp.l3vpn.0 on page 729 show route table bgp.l3vpn.0 detail on page 729 show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 731 show route table inet.0 on page 731 show route table inet6.0 on page 731 show route table inet6.3 on page 732 show route table inetflow detail on page 732 show route table l2circuit.0 on page 732 show route table mpls on page 733 show route table mpls extensive on page 733 show route table mpls.0 on page 733 show route table mpls.0 (RSVP Route—Transit LSP) on page 734 show route table vpls_1 detail on page 734 show route table vpn-a on page 734 show route table vpn-a.mdt.0 on page 735 show route table VPN-A detail on page 735 show route table VPN-AB.inet.0 on page 735 show route table VPN_blue.mvpn-inet6.0 on page 736 show route table VPN-A detail on page 736

[show route table inetflow detail on page 737](#)

Output Fields For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

Sample Output

show route table bgp.l2vpn

```
user@host> show route table bgp.l2vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
    *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.15:100:10.255.71.17/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.15:200:10.255.71.18/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100021, Push 100011(top)
```

show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
    Label operation: Push 182449
    Protocol next hop: 10.255.245.12
    Push 182449
    Indirect next hop: 863a630 297
    State: <Active Int Ext>
    Local AS: 35 Peer AS: 35
    Age: 12:19 Metric2: 1
    Task: BGP_35.10.255.245.12+179
    Announcement bits (1): 0-BGP.0.0.0.0+179
    AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

    Communities: 2914:420 target:11111:1 origin:56:78
    VPN Label: 182449
    Localpref: 100
```

```
Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
    Label operation: Push 182465
    Protocol next hop: 10.255.245.12
    Push 182465
    Indirect next hop: 863a8f0 305
    State: <Active Int Ext>
    Local AS: 35 Peer AS: 35
    Age: 12:19 Metric2: 1
    Task: BGP_35.10.255.245.12+179
    Announcement bits (1): 0-BGP.0.0.0.0+179
  AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
    Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
    VPN Label: 182465
    Localpref: 100
    Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
    Label operation: Push 182465
    Protocol next hop: 10.255.245.12
    Push 182465
    Indirect next hop: 86bd210 330
    State: <Active Int Ext>
    Local AS: 35 Peer AS: 35
    Age: 12:19 Metric2: 1
    Task: BGP_35.10.255.245.12+179
    Announcement bits (1): 0-BGP.0.0.0.0+179
  AS path: 30 10458 14203 2914 11853 11853 11853 11853 6496 6496 6496 6496 6496
    6496 I
    Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
    VPN Label: 182465
    Localpref: 100
    Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
    Label operation: Push 182465
    Protocol next hop: 10.255.245.12
    Push 182465
    Indirect next hop: 86bd210 330
    State: <Active Int Ext>
    Local AS: 35 Peer AS: 35
    Age: 12:19 Metric2: 1
    Task: BGP_35.10.255.245.12+179
    Announcement bits (1): 0-BGP.0.0.0.0+179
  AS path: 30 10458 14203 2914 11853 11853 11853 11853 6496 6496 6496 6496 6496
    6496 I
```



```

Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100

```

show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```

user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
                *[RTarget/5] 00:03:14
                  Type Proxy
                    for 10.255.165.103
                    for 10.255.166.124
                  Local

```

show route table inet.0

```

user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0        *[Static/5] 00:51:57
                  > to 111.222.5.254 via fxp0.0
1.0.0.1/32       *[Direct/0] 00:51:58
                  > via at-5/3/0.0
1.0.0.2/32       *[Local/0] 00:51:58
                  Local
12.12.12.21/32   *[Local/0] 00:51:57
                  Reject
13.13.13.13/32   *[Direct/0] 00:51:58
                  > via t3-5/2/1.0
13.13.13.14/32   *[Local/0] 00:51:58
                  Local
13.13.13.21/32   *[Local/0] 00:51:58
                  Local
13.13.13.22/32   *[Direct/0] 00:33:59
                  > via t3-5/2/0.0
127.0.0.1/32     [Direct/0] 00:51:58
                  > via lo0.0
111.222.5.0/24   *[Direct/0] 00:51:58
                  > via fxp0.0
111.222.5.81/32  *[Local/0] 00:51:58
                  Local

```

show route table inet6.0

```

user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0

```

show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
    *[LDP/9] 00:00:22, metric 1
    > via so-1/0/0.0
::10.255.245.196/128
    *[LDP/9] 00:00:08, metric 1
    > via so-1/0/0.0, Push 100008
```

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP Preference: 170/-101
        Next-hop reference count: 2
        State: **Active Ext>
        Local AS: 65002 Peer AS: 65000
        Age: 4
        Task: BGP_65000.10.12.99.5+3792
        Announcement bits (1): 0-Flow
        AS path: 65000 I
        Communities: traffic-rate:0:0
        Validation state: Accept, Originator: 10.12.99.5
        Via: 10.12.44.0/24, Active
        Localpref: 100
        Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow Preference: 5
        Next-hop reference count: 2
        State: **Active>
        Local AS: 65002
        Age: 6:30
        Task: RT Flow
        Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
        AS path: I
        Communities: 1:1
```

show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    > via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
    *[LDP/9] 00:50:14
    Discard
10.1.1.195:CtrlWord:1:2:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    > via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
```

```
*[LDP/9] 00:50:14
Discard
```

show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 00:13:55, metric 1
                  Receive
1                *[MPLS/0] 00:13:55, metric 1
                  Receive
2                *[MPLS/0] 00:13:55, metric 1
                  Receive
1024             *[VPN/0] 00:04:18
                  to table red.inet.0, Pop
```

show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
    *LDP      Preference: 9
              Next hop: via so-1/0/0.0, selected
              Pop
              State: <Active Int>
              Age: 29:50      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 10.0.0.194/32
```

show route table mpls.0

```
user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 00:45:09, metric 1
                  Receive
1                *[MPLS/0] 00:45:09, metric 1
                  Receive
2                *[MPLS/0] 00:45:09, metric 1
                  Receive
100000           *[L2VPN/7] 00:43:04
                  > via so-0/1/0.1, Pop
100001           *[L2VPN/7] 00:43:03
                  > via so-0/1/0.2, Pop      Offset: 4
100002           *[LDP/9] 00:43:22, metric 1
                  via so-0/1/2.0, Pop
                  > via so-0/1/3.0, Pop
100002(S=0)      *[LDP/9] 00:43:22, metric 1
                  via so-0/1/2.0, Pop
                  > via so-0/1/3.0, Pop
100003           *[LDP/9] 00:43:22, metric 1
                  > via so-0/1/2.0, Swap 100002
                  via so-0/1/3.0, Swap 100002
100004           *[LDP/9] 00:43:16, metric 1
                  via so-0/1/2.0, Swap 100049
                  > via so-0/1/3.0, Swap 100049
```

```

so-0/1/0.1      *[L2VPN/7] 00:43:04
                 > via so-0/1/2.0, Push 100001, Push 100049(top)
                 via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2      *[L2VPN/7] 00:43:03
                 via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
                 > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4

```

show route table mpls.0 (RSVP Route—Transit LSP)

```
user@host> show route table mpls.0
```

```

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

0          *[MPLS/0] 00:37:31, metric 1
           Receive
1          *[MPLS/0] 00:37:31, metric 1
           Receive
2          *[MPLS/0] 00:37:31, metric 1
           Receive
13         *[MPLS/0] 00:37:31, metric 1
           Receive
300352     *[RSVP/7/1] 00:08:00, metric 1
           > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0) *[RSVP/7/1] 00:08:00, metric 1
           > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384     *[RSVP/7/2] 00:05:20, metric 1
           > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0) *[RSVP/7/2] 00:05:20, metric 1
           > to 8.64.1.106 via ge-1/0/0.0, Pop

```

show route table vpls_1 detail

```
user@host> show route table vpls_1 detail
```

```

vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

```

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF

```

show route table vpn-a

```
user@host> show route table vpn-a
```

```
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```

192.168.16.1:1:1/96
           *[VPN/7] 05:48:27
           Discard
192.168.24.1:1:2:1/96
           *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
           AS path: I
           > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

```

192.168.24.1:1:3:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

show route table vpn-a.mdt.0

```

user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
    *[MVPN/70] 01:23:05, metric2 1
    Indirect
1:1:1:10.255.14.218:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
    AS path: I
    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
    AS path: I
    > via so-0/0/1.0, label-switched-path r0-to-r2

```

show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
    *BGP Preference: 170/-101
    Route Distinguisher: 10.255.179.13:200
    Next hop type: Indirect
    Next-hop reference count: 5
    Source: 10.255.179.13
    Next hop type: Router, Next hop index: 732
    Next hop: 10.39.1.14 via fe-0/3/0.0, selected
    Label operation: Push 299824, Push 299824(top)
    Protocol next hop: 10.255.179.13
    Push 299824
    Indirect next hop: 8f275a0 1048574
    State: (Secondary Active Int Ext)
    Local AS: 1 Peer AS: 1
    Age: 3:41:06 Metric: 1 Metric2: 1
    Task: BGP_1.10.255.179.13+64309
    Announcement bits (2): 0-KRT 1-BGP RT Background
    AS path: I
    Communities: target:1:200 rte-type:0.0.0.0:1:0
    Import Accepted
    VPN Label: 299824 TTL Action: vrf-ttl-propagate
    Localpref: 100
    Router ID: 10.255.179.13
    Primary Routing Table bgp.13vpn.0

```

show route table VPN-AB.inet.0

```

user@host> show route table VPN-AB.inet.0
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.0/30      *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0
10.39.1.4/30     *[Direct/0] 00:08:42
                  > via so-5/1/0.0

```

```

10.39.1.6/32      *[Local/0] 00:08:46
                  Local
10.255.71.16/32   *[Static/5] 00:07:24
                  > via so-2/0/0.0
10.255.71.17/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32 *[BGP/170] 00:08:35, localpref 100
                  AS path: 2 I
                  > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32 *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0

```

show route table VPN_blue.mvpn-inet6.0

```

user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
    *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
    AS path: I
    > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
    *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
    AS path: I
    > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
    *[MVPN/70] 00:57:23, metric2 1
    Indirect
5:10.255.2.202:65535:128::192.168.90.2:128:ffff::1/432
    *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
    AS path: I
    > via so-0/1/3.0
6:10.255.2.203:65535:65000:128::10.12.53.12:128:ffff::1/432
    *[PIM/105] 00:02:37
    Multicast (IPv6)
7:10.255.2.202:65535:65000:128::192.168.90.2:128:ffff::1/432
    *[MVPN/70] 00:02:37, metric2 1
    Indirect

```

show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
    *BGP    Preference: 170/-101
            Route Distinguisher: 10.255.179.13:200
            Next hop type: Indirect
            Next-hop reference count: 5
            Source: 10.255.179.13
            Next hop type: Router, Next hop index: 732
            Next hop: 10.39.1.14 via fe-0/3/0.0, selected
            Label operation: Push 299824, Push 299824(top)
            Protocol next hop: 10.255.179.13
            Push 299824

```

```

Indirect next hop: 8f275a0 1048574
State: (Secondary Active Int Ext)
Local AS: 1 Peer AS: 1
Age: 3:41:06 Metric: 1 Metric2: 1
Task: BGP_1.10.255.179.13+64309
Announcement bits (2): 0-KRT 1-BGP RT Background
AS path: I
Communities: target:1:200 rte-type:0.0.0.0:1:0
Import Accepted
VPN Label: 299824 TTL Action: vrf-ttl-propagate
Localpref: 100
Router ID: 10.255.179.13
Primary Routing Table bgp.l3vpn.0

```

show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP    Preference: 170/-101
            Next-hop reference count: 2
            State: **Active Ext>
            Local AS: 65002 Peer AS: 65000
            Age: 4
            Task: BGP_65000.10.12.99.5+3792
            Announcement bits (1): 0-Flow
            AS path: 65000 I
            Communities: traffic-rate:0:0
            Validation state: Accept, Originator: 10.12.99.5
            Via: 10.12.44.0/24, Active
            Localpref: 100
            Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow    Preference: 5
            Next-hop reference count: 2
            State: **Active>
            Local AS: 65002
            Age: 6:30
            Task: RT Flow
            Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
            AS path: I
            Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1.1.1.4/96 AD
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/1.5
1.1.1.2:100:1.0/96 MH
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1.0/96 MH
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified

```

```

> via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
*[VPLS/7] 1d 03:11:02, metric2 1
> via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
*[LDP/9] 1d 03:11:02
Discard
```


show route terse


Syntax	show route terse <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route terse
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display a high-level summary of the routes in the routing table.
<div>  <p>NOTE: For BGP routes, the show route terse command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.</p> <p>To display the metric1 and metric2 value of a BGP route, use the show route extensive command.</p> </div>	
Options	<p>none—Display a high-level summary of the routes in the routing table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route terse on page 741
Output Fields	Table 73 on page 739 describes the output fields for the show route terse command. Output fields are listed in the approximate order in which they appear.

Table 73: show route terse Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active) holddown (routes that are in the pending state before being declared inactive) hidden (routes that are not used because of a routing policy)

Table 73: show route terse Output Fields (*continued*)

Field Name	Field Description
route key	Key for the state of the route: <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.
A	Active route. An asterisk (*) indicates this is the active route.
V	Validation status of the route: <ul style="list-style-type: none"> • ?—Not evaluated. Indicates that the route was not learned through BGP. • I—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • N—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database. • V—Valid. Indicates that the prefix and autonomous system pair are found in the database.
Destination	Destination of the route.
P	Protocol through which the route was learned: <ul style="list-style-type: none"> • A—Aggregate • B—BGP • C—CCC • D—Direct • G—GMPLS • I—IS-IS • L—L2CKT, L2VPN, LDP, Local • K—Kernel • M—MPLS, MSDP • O—OSPF • P—PIM • R—RIP, RIPng • S—Static • T—Tunnel
Prf	Preference value of the route. In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.
Metric 1	First metric value in the route. For routes learned from BGP, this is the MED metric.
Metric 2	Second metric value in the route. For routes learned from BGP, this is the IGP metric.

Table 73: show route terse Output Fields (*continued*)

Field Name	Field Description
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> I—IGP. E—EGP. ?—Incomplete; typically, the AS path was aggregated.

Sample Output

show route terse

```

user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* ? 1.0.1.1/32        0 10      1           >10.0.0.2
?                               B 170      100           I
  unverified
* ? 1.1.1.1/32        D 0           >10.0.0.2
* ? 1.1.1.1/32        D 0           >10.0.0.2
* V 2.2.0.2/32        B 170     110         >10.0.0.2      200 I
  valid
* ? 10.0.0.0/30       D 0           >10.0.0.2
* ? 10.0.0.0/30       B 170     100         >1t-1/2/0.1    I
  unverified
* ? 10.0.0.1/32       L 0           Local
* ? 10.0.0.4/30       B 170     100         >10.0.0.2      I
  unverified
* ? 10.0.0.8/30       B 170     100         >10.0.0.2      I
  unverified
* I 172.16.1.1/32     B 170      90         >10.0.0.2      200 I
  invalid
* N 192.168.2.3/32    B 170     100         >10.0.0.2      200 I
  unknown
* ? 224.0.0.5/32      O 10      1           MultiRecv

```

