



Junos[®] OS

IS-IS Feature Guide for Routing Devices

Release

14.1



Published: 2014-05-13

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

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

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

Junos[®] OS IS-IS Feature Guide for Routing Devices

14.1

Copyright © 2014, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	xv
	Documentation and Release Notes	xv
	Supported Platforms	xv
	Using the Examples in This Manual	xvi
	Merging a Full Example	xvi
	Merging a Snippet	xvii
	Documentation Conventions	xvii
	Documentation Feedback	xix
	Requesting Technical Support	xx
	Self-Help Online Tools and Resources	xx
	Opening a Case with JTAC	xx
Part 1	Overview	
Chapter 1	Introduction to IS-IS	3
	IS-IS Overview	3
	IS-IS Terminology	3
	ISO Network Addresses	4
	IS-IS Packets	6
	Persistent Route Reachability	7
	IS-IS Support for Multipoint Network Clouds	7
	Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels	7
Chapter 2	IS-IS Standards	9
	Supported IS-IS Standards	9
Part 2	Configuration	
Chapter 3	Basic IS-IS Configuration	13
	Example: Configuring IS-IS	13
	Understanding IS-IS Configuration	13
	Example: Configuring IS-IS	14
	Example: Configuring Multi-Level IS-IS	19
	Understanding IS-IS Areas	19
	Example: Configuring Multi-Level IS-IS	20
	Example: Configuring IS-IS on Logical Systems Within the Same Router	28
	Introduction to Logical Systems	28
	Example: Configuring IS-IS on Logical Systems Within the Same Router	30

	Example: Configuring IS-IS Designated Routers	39
	Understanding IS-IS Designated Routers	39
	Example: Configuring Designated Router Election Priority for IS-IS	40
	Configuring an ISO System Identifier for the Router	40
	Example: Configuring the Name of the Router, IP Address, and System ID	41
Chapter 4	IS-IS Authentication and Checksums	43
	Configuring IS-IS Authentication	43
	Configuring Authentication Without Network-Wide Deployment	45
	Example: Configuring Hitless Authentication Key Rollover for IS-IS	45
	Understanding Hitless Authentication Key Rollover for IS-IS	45
	Example: Configuring Hitless Authentication Key Rollover for IS-IS	46
	Example: Enabling Packet Checksums on IS-IS Interfaces	50
	Understanding Checksums on IS-IS Interfaces	50
	Example: Enabling Packet Checksums on IS-IS Interfaces	51
Chapter 5	IS-IS Routing Policy and Route Redistribution	55
	Example: Configuring IS-IS Route Leaking From a Level 2 Area to a Level 1 Area	55
	Example: Redistributing OSPF Routes into IS-IS	61
	Understanding Routing Policies	61
	Importing and Exporting Routes	62
	Active and Inactive Routes	63
	Explicitly Configured Routes	63
	Dynamic Database	64
	Example: Redistributing OSPF Routes into IS-IS	64
	Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS	72
	Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions	72
	Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS	73
	Example: Configuring an IS-IS Default Route Policy on Logical Systems	82
	Understanding Default Routes	82
	Example: Configuring an IS-IS Default Route Policy on Logical Systems	82
	IS-IS Extensions to Support Route Tagging	88
Chapter 6	IS-IS Bidirectional Forwarding Detection	91
	Example: Configuring BFD for IS-IS	91
	Understanding BFD for IS-IS	91
	Example: Configuring BFD for IS-IS	93
	Example: Configuring BFD Authentication for IS-IS	99
	Understanding BFD Authentication for IS-IS	99
	BFD Authentication Algorithms	100
	Security Authentication Keychains	100
	Strict Versus Loose Authentication	101
	Configuring BFD Authentication for IS-IS	101
	Configuring BFD Authentication Parameters	101
	Viewing Authentication Information for BFD Sessions	102
	Example: Configuring BFD Authentication for IS-IS	104

Chapter 7	IS-IS Multitopology Routing and IPv6 Support	109
	Example: Configuring IS-IS Multicast Topology	109
	IS-IS Multicast Topologies Overview	109
	Example: Configuring IS-IS Multicast Topology	110
	Example: Configuring IS-IS Dual Stacking of IPv4 and IPv6 Unicast Addresses	124
	Understanding Dual Stacking	125
	Example: Configuring IS-IS Dual Stacking of IPv4 and IPv6 Unicast Addresses	126
	Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies	131
	Understanding IS-IS IPv4 and IPv6 Unicast Topologies	132
	Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies	132
Chapter 8	IS-IS Link and Node Protection	141
	Example: Configuring Link and Node Protection for IS-IS Routes	141
	Understanding Loop-Free Alternate Routes for IS-IS	141
	Configuring Link Protection for IS-IS	143
	Configuring Node-Link Protection for IS-IS	144
	Excluding an IS-IS Interface as a Backup for Protected Interfaces	144
	Configuring RSVP Label-Switched Paths as Backup Paths for IS-IS	144
	Using Operational Mode Commands to Monitor Protected IS-IS Routes	145
	Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN	145
Chapter 9	IS-IS Traffic Engineering	157
	IS-IS Extensions to Support Traffic Engineering	157
	IS-IS IGP Shortcuts	157
	Example: Enabling IS-IS Traffic Engineering Support	158
	IGP Shortcuts	158
	Example: Enabling IS-IS Traffic Engineering Support	159
	Example: Advertising Label-Switched Paths into IS-IS	172
	Understanding Forwarding Adjacencies	172
	Example: Advertising Label-Switched Paths into IS-IS	173
	Example: Enabling Wide IS-IS Metrics for Traffic Engineering	180
	Understanding Wide IS-IS Metrics for Traffic Engineering	180
	Example: Enabling Wide IS-IS Metrics for Traffic Engineering	181
	Example: Configuring Synchronization Between IS-IS and LDP	183
	Understanding LDP-IGP Synchronization	183
	Synchronization Behavior During Graceful Restart	184
	Synchronization Behavior on LAN Interfaces	184
	Synchronization Behavior on IGP Passive Interfaces	184
	Synchronization and TE Metrics	185
	Example: Configuring Synchronization Between IS-IS and LDP	185

Chapter 10	IS-IS Scaling and Throttling	189
	Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces	189
	Understanding Link-State PDU Throttling for IS-IS Interfaces	189
	Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces	190
	Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces	194
	Understanding the Transmission Frequency for CSNPs on IS-IS Interfaces	194
	Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces	195
	Example: Configuring Mesh Groups of IS-IS Interfaces	199
	Understanding IS-IS Mesh Groups	199
	Example: Configuring Mesh Groups of IS-IS Interfaces	200
Chapter 11	IS-IS CLNS	207
	Example: Configuring IS-IS for CLNS	207
	Understanding IS-IS for CLNS	207
	Example: Configuring IS-IS for CLNS	207
Chapter 12	Configuration Statements	211
	[edit protocols isis] Hierarchy Level	213
	authentication-key (Protocols IS-IS)	216
	authentication-key-chain (Protocols IS-IS)	217
	authentication-type (Protocols IS-IS)	218
	backup-spf-options (Protocols IS-IS)	219
	bfd-liveness-detection (Protocols IS-IS)	220
	checksum (Protocols IS-IS)	222
	clns-routing	222
	clns-updown-compatibility	223
	context-identifier (Protocols IS-IS)	224
	csnp-interval	225
	disable (Protocols IS-IS)	226
	disable (LDP Synchronization for IS-IS)	227
	export (Protocols IS-IS)	228
	external-preference (Protocols IS-IS)	229
	family (Protocols IS-IS)	230
	graceful-restart (Protocols IS-IS)	231
	hello-authentication-key	232
	hello-authentication-key-chain	233
	hello-authentication-type	234
	hello-interval (Protocols IS-IS)	235
	hello-padding	236
	hold-time (Protocols IS-IS)	238
	hold-time (LDP Synchronization for IS-IS)	239
	ignore-attached-bit	240
	ignore-lsp-metrics (Protocols IS-IS)	241
	interface (Protocols IS-IS)	242

ipv4-multicast	244
ipv4-multicast-metric	245
ipv6-multicast	245
ipv6-multicast-metric	246
ipv6-unicast	247
ipv6-unicast-metric	248
isis	249
label-switched-path (Protocols IS-IS)	250
ldp-synchronization	251
level (Global IS-IS)	252
level (IS-IS Interfaces)	253
link-protection (Protocols IS-IS)	254
loose-authentication-check	254
lsp-equal-cost	255
lsp-interval	256
lsp-lifetime	257
max-areas	258
max-hello-size	259
max-lsp-size	259
max-snp-size	260
mesh-group (Protocols IS-IS)	261
metric (Protocols IS-IS)	262
multicast-rpf-routes	263
multipath (Protocols IS-IS)	264
no-adjacency-down-notification (Protocols IS-IS)	265
no-adjacency-holddown	266
no-authentication-check	266
no-csnp-authentication	267
no-eligible-backup (Protocols IS-IS)	267
no-hello-authentication	268
no-ipv4-multicast	268
no-ipv4-routing	269
no-ipv6-multicast	270
no-ipv6-routing	271
no-ipv6-unicast	272
no-psnp-authentication	272
no-unicast-topology	273
node-link-protection (Protocols IS-IS)	273
overload (Protocols IS-IS)	274
passive (Protocols IS-IS)	277
point-to-point	278
preference (Protocols IS-IS)	279
prefix-export-limit (Protocols IS-IS)	280
priority (Protocols IS-IS)	281
protocols	282
reference-bandwidth (Protocols IS-IS)	284
rib-group (Protocols IS-IS)	285
routing-instances (Multiple Routing Entities)	286
shortcuts (Protocols IS-IS)	287

spf-options (Protocols IS-IS)	288
static-host-mapping	290
te-metric (Protocols IS-IS)	291
topologies (Protocols IS-IS)	292
traceoptions (Protocols IS-IS)	293
traffic-engineering (Protocols IS-IS)	296
wide-metrics-only	299

Part 3

Chapter 13

Administration

Operational Commands 303

clear bfd adaptation	305
clear bfd session	306
clear isis adjacency	307
clear isis database	309
clear isis overload	311
clear isis statistics	313
ping clns	315
restart	317
show bfd session	327
show isis adjacency	335
show isis authentication	339
show isis backup coverage	341
show isis backup label-switched-path	343
show isis backup spf results	345
show isis context-identifier	349
show isis database	351
show isis hostname	363
show isis interface	365
show isis overview	369
show isis route	372
show isis spf	376
show isis statistics	381
show policy	384
show policy conditions	386
show route	388
show route active-path	393
show route advertising-protocol	398
show route all	403
show route best	405
show route brief	408
show route detail	410
show route exact	427
show route export	429
show route extensive	432
show route forwarding-table	449
show route hidden	463
show route inactive-path	466
show route instance	469

	show route next-hop	476
	show route output	482
	show route protocol	487
	show route receive-protocol	499
	show route table	507
	show route terse	521
	show security keychain	524
	test policy	527
	traceroute clns	529
Part 4	Troubleshooting	
Chapter 14	IS-IS Troubleshooting	533
	Verifying the IS-IS Protocol	533
	Verify the LSP	534
	Verify IS-IS Adjacencies and Interfaces	535
	Verify the IS-IS Configuration	536
	Take Appropriate Action	537
	Verify the LSP Again	538
	Verifying the IS-IS Configuration on a Router in a Network	541
	Check the Configuration of a Level 1/Level 2 Router	542
	Check the Configuration of a Level 1 Router	544
	Check the Configuration of a Level 2 Router	545
	Displaying the Status of IS-IS Adjacencies	547
	Displaying Detailed IS-IS Protocol Information	548
	Analyzing IS-IS Link-State PDUs in Detail	550
	Displaying Sent or Received IS-IS Protocol Packets	552
Chapter 15	Routing Protocol Process Memory FAQs	555
	Routing Protocol Process Memory FAQs Overview	555
	Routing Protocol Process Memory FAQs	556
	Frequently Asked Questions: Routing Protocol Process Memory	556
	Frequently Asked Questions: Interpreting Routing Protocol Process-Related Command Outputs	557
	Frequently Asked Questions: Routing Protocol Process Memory Swapping	560
	Frequently Asked Questions: Troubleshooting the Routing Protocol Process	561
Part 5	Index	
	Index	565

List of Figures

Part 1	Overview	
Chapter 1	Introduction to IS-IS	3
	Figure 1: Install Default Route to Nearest Routing Device That Operates at Both Level 1 and Level 2	7
Part 2	Configuration	
Chapter 3	Basic IS-IS Configuration	13
	Figure 2: Simple IS-IS Topology	15
	Figure 3: IS-IS Multi-Level Topology	21
	Figure 4: Logical Systems Concepts	29
	Figure 5: Junos OS Without Logical Systems	30
	Figure 6: Junos OS With Logical Systems	30
	Figure 7: IS-IS on Logical Systems	31
Chapter 4	IS-IS Authentication and Checksums	43
	Figure 8: Hitless Authentication Key Rollover for IS-IS	47
	Figure 9: IS-IS Checksum Topology	51
Chapter 5	IS-IS Routing Policy and Route Redistribution	55
	Figure 10: Route Leaking From a Level 2 Area to a Level 1 Area	56
	Figure 11: Importing and Exporting Routes	62
	Figure 12: IS-IS Route Redistribution Topology	65
	Figure 13: Redistributing BGP Routes with a Specific Community Tag into IS-IS	74
	Figure 14: IS-IS Logical Systems with a Default Route to an ISP	83
Chapter 6	IS-IS Bidirectional Forwarding Detection	91
	Figure 15: Configuring BFD for IS-IS	94
	Figure 16: IS-IS BFD Authentication Topology	104
Chapter 7	IS-IS Multitopology Routing and IPv6 Support	109
	Figure 17: Configuring IS-IS Multicast Topology	111
	Figure 18: IS-IS IPv4 and IPv6 Dual Stacking Topology	127
	Figure 19: IS-IS IPv4 and IPv6 Unicast Topologies	134
Chapter 8	IS-IS Link and Node Protection	141
	Figure 20: Link Protection and Node-Link Protection Comparison for IS-IS Routes	143
	Figure 21: IS-IS Node-Link Protection Topology	146
Chapter 9	IS-IS Traffic Engineering	157

	Figure 22: Typical SPF Tree, Sourced from Router A	158
	Figure 23: Modified SPF Tree, Using LSP A–D as a Shortcut	158
	Figure 24: Modified SPF Tree, Using LSP A–D and LSP A–E as Shortcuts	159
	Figure 25: IS-IS Shortcuts Topology	160
	Figure 26: IS-IS Advertising a Label-Switched Path Topology	173
	Figure 27: IS-IS Wide Metrics Topology	181
	Figure 28: IS-IS and LDP Synchronization Topology	186
Chapter 10	IS-IS Scaling and Throttling	189
	Figure 29: IS-IS Link-State PDU Interval Topology	190
	Figure 30: IS-IS CSNP Interval Topology	195
	Figure 31: IS-IS Mesh Topology	200
Part 4	Troubleshooting	
Chapter 14	IS-IS Troubleshooting	533
	Figure 32: MPLS Network Broken at the IS-IS Protocol Layer	533
	Figure 33: Levels in an IS-IS Network Topology	541
	Figure 34: IS-IS Network Topology with Details	542
	Figure 35: IS-IS Network Topology	548

List of Tables

	About the Documentation	xv
	Table 1: Notice Icons	xviii
	Table 2: Text and Syntax Conventions	xviii
Part 2	Configuration	
Chapter 6	IS-IS Bidirectional Forwarding Detection	91
	Table 3: Configuring BFD for IS-IS	92
Chapter 7	IS-IS Multitopology Routing and IPv6 Support	109
	Table 4: IPv4 Statements	110
	Table 5: IPv6 Statements	110
	Table 6: Choosing the Right Solution to Address Next-Generation Addressing Requirements	125
Chapter 12	Configuration Statements	211
	Table 7: Default Metric Values for Routes Exported into IS-IS	262
Part 3	Administration	
Chapter 13	Operational Commands	303
	Table 8: show bfd session Output Fields	328
	Table 9: show isis adjacency Output Fields	335
	Table 10: show isis authentication Output Fields	339
	Table 11: show isis backup coverage Output Fields	341
	Table 12: show isis backup label-switched-path Output Fields	343
	Table 13: show isis backup spf results Output Fields	346
	Table 14: show isis context-identifier Output Fields	349
	Table 15: show isis database Output Fields	352
	Table 16: show isis hostname Output Fields	363
	Table 17: show isis interface Output Fields	366
	Table 18: show isis overview Output Fields	369
	Table 19: show isis route Output Fields	372
	Table 20: show isis spf Output Fields	376
	Table 21: show isis statistics Output Fields	382
	Table 22: show policy Output Fields	384
	Table 23: show policy conditions Output Fields	386
	Table 24: show route Output Fields	389
	Table 25: show route advertising-protocol Output Fields	399
	Table 26: show route detail Output Fields	410
	Table 27: Next-hop Types Output Field Values	415
	Table 28: State Output Field Values	416

	Table 29: Communities Output Field Values	418
	Table 30: show route export Output Fields	429
	Table 31: show route extensive Output Fields	432
	Table 32: show route forwarding-table Output Fields	452
	Table 33: show route instance Output Fields	470
	Table 34: show route receive-protocol Output Fields	499
	Table 35: show route terse Output Fields	521
	Table 36: show security keychain Output Fields	524
	Table 37: Traceroute clns Output Fields	529
Part 4	Troubleshooting	
Chapter 14	IS-IS Troubleshooting	533
	Table 38: IS-IS Protocol Tracing Flags	549
Chapter 15	Routing Protocol Process Memory FAQs	555
	Table 39: show system processes extensive Output Fields	558
	Table 40: show task memory Output Fields	559

About the Documentation

- [Documentation and Release Notes on page xv](#)
- [Supported Platforms on page xv](#)
- [Using the Examples in This Manual on page xvi](#)
- [Documentation Conventions on page xvii](#)
- [Documentation Feedback on page xix](#)
- [Requesting Technical Support on page xx](#)

Documentation and Release Notes

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

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

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

Supported Platforms

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

- [ACX Series](#)
- [J Series](#)
- [T Series](#)
- [MX Series](#)
- [SRX Series](#)
- [M Series](#)

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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


Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

[Table 1 on page xviii](#) defines notice icons used in this guide.

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>

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

Convention	Description	Examples
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <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)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	<pre>[edit] routing-options { static { route default { nexthop <i>address</i>; retain; } } }</pre>
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number

- Software release version (if applicable)

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [Introduction to IS-IS on page 3](#)
- [IS-IS Standards on page 9](#)

CHAPTER 1

Introduction to IS-IS

- [IS-IS Overview on page 3](#)

IS-IS Overview

The IS-IS protocol is an interior gateway protocol (IGP) that uses link-state information to make routing decisions.

IS-IS is a link-state IGP that uses the shortest-path-first (SPF) algorithm to determine routes. IS-IS evaluates the topology changes and determines whether to perform a full SPF recalculation or a partial route calculation (PRC). This protocol originally was developed for routing International Organization for Standardization (ISO) Connectionless Network Protocol (CLNP) packets.

Like OSPF routing, IS-IS uses hello packets that allow network convergence to occur quickly when network changes are detected. IS-IS uses the SPF algorithm to determine routes. Using SPF, IS-IS evaluates network topology changes and determines if a full or partial route calculation is required.



NOTE: Because IS-IS uses ISO addresses, the configuration of IP version 6 (IPv6) and IP version 4 (IPv4) implementations of IS-IS is identical.

This section discusses the following topics:

- [IS-IS Terminology on page 3](#)
- [ISO Network Addresses on page 4](#)
- [IS-IS Packets on page 6](#)
- [Persistent Route Reachability on page 7](#)
- [IS-IS Support for Multipoint Network Clouds on page 7](#)
- [Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels on page 7](#)

IS-IS Terminology

An IS-IS network is a single autonomous system (AS), also called a *routing domain*, that consists of *end systems* and *intermediate systems*. End systems are network entities that

send and receive packets. Intermediate systems send and receive packets and relay (forward) packets. (Intermediate system is the Open System Interconnection [OSI] term for a router.) ISO packets are called network PDUs.

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. No IS-IS area functions strictly as a backbone.

Level 1 routers share intra-area routing information, and Level 2 routers share interarea information about IP addresses available within each area. Uniquely, IS-IS routers can act as both Level 1 and Level 2 routers, sharing intra-area routes with other Level 1 routers and interarea routes with other Level 2 routers.

The propagation of link-state updates is determined by the level boundaries. All routers within a level maintain a complete link-state database of all other routers in the same level. Each router then uses the Dijkstra algorithm to determine the shortest path from the local router to other routers in the link-state database.

ISO Network Addresses

IS-IS uses ISO network addresses. Each address identifies a point of connection to the network, such as a router interface, and is called a *network service access point (NSAP)*.

IS-IS supports multiple NSAP addresses on the loopback lo0 interface.

An end system can have multiple NSAP addresses, in which case the addresses differ only by the last byte (called the *n-selector*). Each NSAP represents a service that is available at that node. In addition to having multiple services, a single node can belong to multiple areas.

Each network entity also has a special network address called a *network entity title (NET)*. Structurally, an NET is identical to an NSAP address but has an n-selector of 00. Most end systems and intermediate systems have one NET. Intermediate systems that participate in multiple areas can have multiple NETs.

The following ISO addresses illustrate the IS-IS address format:

```
49.0001.00a0.c96b.c490.00
49.0001.2081.9716.9018.00
```

NETs take several forms, depending on your network requirements. NET addresses are hexadecimal and range from 8 octets to 20 octets in length. Generally, the format consists of an authority and format Identifier (AFI), a domain ID, an area ID, a system identifier, and a selector. The simplest format omits the domain ID and is 10 octets long. For example, the NET address 49.0001.1921.6800.1001.00 consists of the following parts:

- 49—AFI
- 0001—Area ID

- 1921.6800.1001—System identifier
- 00—Selector

The system identifier must be unique within the network. For an IP-only network, we recommend using the IP address of an interface on the router. Configuring a loopback NET address with the IP address is helpful when troubleshooting is required on the network.

The first portion of the address is the area number, which is a variable number from 1 through 13 bytes. The first byte of the area number (49) is the authority and format indicator (AFI). The next bytes are the assigned domain (area) identifier, which can be from 0 through 12 bytes. In the examples above, the area identifier is 0001.

The next six bytes form the system identifier. The system identifier can be any six bytes that are unique throughout the entire domain. The system identifier commonly is the media access control (MAC) address (as in the first example, 00a0.c96b.c490) or the IP address expressed in binary-coded decimal (BCD) (as in the second example, 2081.9716.9018, which corresponds to IP address 208.197.169.18). The last byte (00) is the n-selector.



NOTE: The system identifier cannot be 0000.0000.0000. All 0s is an illegal setting, and the adjacency is not formed with this setting.

To provide help with IS-IS debugging, the Junos[®] operating system (Junos OS) supports dynamic mapping of ISO system identifiers to the hostname. Each system can be configured with a hostname, which allows the system identifier-to-hostname mapping to be carried in a dynamic hostname type, length, and value (TLV) tuple in IS-IS link-state PDUs. This enables intermediate systems in the routing domain to learn about the ISO system identifier of a particular intermediate system.

IS-IS Packets

Each IS-IS PDU shares a common header. IS-IS uses the following PDUs to exchange protocol information:

- IS-IS hello (IIH) PDUs—Broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems.

IS-IS hello PDUs establish adjacencies with other routers and have three different formats: one for point-to-point hello packets, one for Level 1 broadcast links, and one for Level 2 broadcast links. Level 1 routers must share the same area address to form an adjacency, while Level 2 routers do not have this limitation. The request for adjacency is encoded in the Circuit type field of the PDU.

Hello PDUs have a preset length assigned to them. The IS-IS router does not resize any PDU to match the maximum transmission unit (MTU) on a router interface. Each interface supports the maximum IS-IS PDU of 1492 bytes, and hello PDUs are padded to meet the maximum value. When the hello is sent to a neighboring router, the connecting interface supports the maximum PDU size.

- Link-state PDUs—Contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area.

Also included is metric and IS-IS neighbor information. Each link-state PDU must be refreshed periodically on the network and is acknowledged by information within a sequence number PDU.

On point-to-point links, each link-state PDU is acknowledged by a partial sequence number PDU (PSNP), but on broadcast links, a complete sequence number PDU (CSNP) is sent out over the network. Any router that finds newer link-state PDU information in the CSNP then purges the out-of-date entry and updates the link-state database.

Link-state PDUs support variable-length subnet mask addressing.

- Complete sequence number PDUs (CSNPs)—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.

Contained within the CSNP is a link-state PDU identifier, a lifetime, a sequence number, and a checksum for each entry in the database. Periodically, a CSNP is sent on both broadcast and point-to-point links to maintain a correct database. Also, the advertisement of CSNPs occurs when an adjacency is formed with another router. Like IS-IS hello PDUs, CSNPs come in two types: Level 1 and Level 2.

When a device receives a CSNP, it checks the database entries against its own local link-state database. If it detects missing information, the device requests specific link-state PDU details using a partial sequence number PDU (PSNP).

- Partial sequence number PDUs (PSNPs)—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.

A PSNP is used by an IS-IS router to request link-state PDU information from a neighboring router. A PSNP can also explicitly acknowledge the receipt of a link-state PDU on a point-to-point link. On a broadcast link, a CSNP is used as implicit knowledge. Like hello PDUs and CSNPs, the PSNP also has two types: Level 1 and Level 2.

When a device compares a CSNP to its local database and determines that a link-state PDU is missing, the router issues a PSNP for the missing link-state PDU, which is returned in a link-state PDU from the router sending the CSNP. The received link-state PDU is then stored in the local database, and an acknowledgment is sent back to the originating router.

Persistent Route Reachability

IPv4 and IPv6 route reachability information in IS-IS link-state PDUs is preserved when you commit a configuration. IP prefixes are preserved with their original packet fragment upon link-state PDU regeneration.


IS-IS Support for Multipoint Network Clouds

IS-IS does not support multipoint configurations. Therefore, when configuring Frame Relay or Asynchronous Transfer Mode (ATM) networks, you must configure them as collections of point-to-point links, not as multipoint clouds.

Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels

When a routing device that operates as both a Level 1 and Level 2 router (Router B) determines that it can reach at least one area other than its own (for example, in Area Y), it sets the ATTACHED bit in its Level 1 link-state PDU. Thereafter, the Level 1 router (Router A) introduces a default route pointing to the nearest attached routing device that operates as both a Level 1 and Level 2 router (Router B). See [Figure 1 on page 7](#).

Figure 1: Install Default Route to Nearest Routing Device That Operates at Both Level 1 and Level 2



Related Documentation

- *IS-IS Feature Guide for Routing Devices*

CHAPTER 2

IS-IS Standards

- [Supported IS-IS Standards on page 9](#)

Supported IS-IS Standards

Junos OS substantially supports the following standards for IS-IS.

- International Organization for Standardization/International Electrotechnical Commission (ISO/IEC) 8473, *Information technology — Protocol for providing the connectionless-mode network service*
- ISO 9542, *End System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for the Provision of the Connectionless-mode Network Service*
- ISO/IEC 10589, *Information technology — Telecommunications and information exchange between systems — Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)*
- RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*
- RFC 3719, *Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)*
- RFC 3847, *Restart Signaling for Intermediate System to Intermediate System (IS-IS)*
- RFC 5120, *M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)*
- RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*
- RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*
- RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*
- RFC 5302, *Domain-Wide Prefix Distribution with Two-Level IS-IS*
- RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*
- RFC 5304, *IS-IS Cryptographic Authentication*
- RFC 5305, *IS-IS Extensions for Traffic Engineering*
- RFC 5306, *Restart Signaling for IS-IS*

- RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*
- RFC 5308, *Routing IPv6 with IS-IS*
- RFC 5310, *IS-IS Generic Cryptographic Authentication*
- RFC 5880, *Bidirectional Forwarding Detection (BFD)*

The following RFCs do not define standards, but provide information about IS-IS and related technologies. The IETF classifies them as “Informational.”

- RFC 2973, *IS-IS Mesh Groups*
- RFC 3358, *Optional Checksums in Intermediate System to Intermediate System (ISIS)*
- RFC 3359, *Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System*
- RFC 3373, *Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies*
- RFC 3567, *Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication*
- RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)*
- RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
- Internet draft draft-ietf-isis-wg-255adj-02.txt, *Maintaining more than 255 circuits in IS-IS*

**Related
Documentation**

- [IS-IS Overview on page 3](#)
- *Supported ES-IS Standards*
- *Accessing Standards Documents on the Internet*

PART 2

Configuration

- [Basic IS-IS Configuration on page 13](#)
- [IS-IS Authentication and Checksums on page 43](#)
- [IS-IS Routing Policy and Route Redistribution on page 55](#)
- [IS-IS Bidirectional Forwarding Detection on page 91](#)
- [IS-IS Multitopology Routing and IPv6 Support on page 109](#)
- [IS-IS Link and Node Protection on page 141](#)
- [IS-IS Traffic Engineering on page 157](#)
- [IS-IS Scaling and Throttling on page 189](#)
- [IS-IS CLNS on page 207](#)
- [Configuration Statements on page 211](#)

CHAPTER 3

Basic IS-IS Configuration

- [Example: Configuring IS-IS on page 13](#)
- [Example: Configuring Multi-Level IS-IS on page 19](#)
- [Example: Configuring IS-IS on Logical Systems Within the Same Router on page 28](#)
- [Example: Configuring IS-IS Designated Routers on page 39](#)
- [Configuring an ISO System Identifier for the Router on page 40](#)
- [Example: Configuring the Name of the Router, IP Address, and System ID on page 41](#)

Example: Configuring IS-IS

- [Understanding IS-IS Configuration on page 13](#)
- [Example: Configuring IS-IS on page 14](#)

Understanding IS-IS Configuration

Configuring IS-IS is straightforward.

To configure IS-IS, you must enable IS-IS on the interfaces and configure a NET address on one of the device interfaces (preferably, the lo0 interface) by setting **family iso address net-address** on the interface. To create the NET address (also known as the system ID or the NSAP address), you can use the convention that is dictated by your network design, or you can follow this convention:

1. Take the router ID, remove the dots (.), and insert leading zeroes where necessary so that the string is 12 characters long.

For example, if the router ID is 192.168.0.4, the 12-character string would be 192168000004. If the router ID is 10.12.23.1, the 12-character string would be 010012023001.

2. Add a dot after every 4th character.

The strings would become 1921.6800.0004 and 0100.1202.3001.

3. Prepend the area number.

If the routing devices are in area 47, the strings would become 47.1921.6800.0004 and 47.0100.1202.3001.

4. Append the selector (00).

The strings would become 47.1921.6800.0004.00 and 47.0100.1202.3001.00.

You must configure the ISO family on all interfaces that are supporting the IS-IS protocol by setting **family iso** on the interface. This means that IS-IS related frames are not discarded by the routing devices.

You must enable IS-IS to run on the interfaces by setting **interface interface-name** in the protocol configuration. This means that the interfaces are advertised into IS-IS.

Unlike OSPF, when you enable IS-IS on the lo0 interface, you do not need to explicitly set passive mode. Passive mode means that the interface is advertised into the link-state protocol, but the interface does not send or receive protocol control packets, such as IS-IS hello and link-state PDUs. In IS-IS, the lo0 interface is always passive.

When you enable IS-IS on an interface, both levels (Level 1 and Level 2) are enabled by default. To specify that an interface is on a Level 1 link, disable Level 2. To specify that an interface is on a Level 2 link, disable Level 1. You can disable a level on the entire device or per-interface. If two routing devices, R1 and R2, are both in the same IS-IS area, they communicate at Level 1 if one or both devices have Level 2 disabled.

For security devices only, you must enable IS-IS by setting **mode packet-based** at the **[edit security forwarding-options family iso]** hierarchy level.

Example: Configuring IS-IS

This example shows how to configure IS-IS.

- [Requirements on page 14](#)
- [Overview on page 14](#)
- [Configuration on page 15](#)
- [Verification on page 17](#)

Requirements

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


Overview

In this example, you configure the two IS-IS routing devices in a single area. The devices have NET addresses 49.0002.0192.0168.0001.00 and 49.0002.0192.0168.0002.00 on the lo0 interfaces. Additionally, you configure the ISO family on the IS-IS interfaces.

For Junos OS security devices only, you configure the **mode packet-based** statement at the **[edit security forwarding-options family iso]** hierarchy level.

[Figure 2 on page 15](#) shows the topology used in this example.

Figure 2: Simple IS-IS Topology



“CLI Quick Configuration” on page 15 shows the configuration for both of the devices in Figure 2 on page 15. The section “Step-by-Step Procedure” on page 15 describes the steps on Device R1.

Configuration

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

Device R1

```

set security forwarding-options family iso mode packet-based
set interfaces ge-1/2/0 unit 0 description to-R2
set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set protocols isis interface ge-1/2/0.0
set protocols isis interface lo0.0
  
```

Device R2

```

set security forwarding-options family iso mode packet-based
set interfaces ge-1/2/0 unit 0 description to-R1
set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
set protocols isis interface ge-1/2/0.0
set protocols isis interface lo0.0
  
```

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

To configure IS-IS:

1. Enable IS-IS if your router is in secure context.


```

[edit security forwarding-options family iso]
user@R1# set mode packet-based
      
```
2. Create the interface that connects to Device R2, and configure the ISO family on the interface.


```

[edit interfaces ge-1/2/0 unit 0]
user@R1# set description to-R2
user@R1# set family inet address 10.0.0.1/30
user@R1# set family iso
      
```

3. Create the loopback interface, set the IP address, and set the NET address.

```
[edit interfaces lo0 unit 0]
user@R1# set family inet address 192.168.0.1/32
user@R1# set family iso address 49.0002.0192.0168.0001.00
```

4. Enable IS-IS on the interfaces.

```
[edit protocols isis]
user@R1# set interface ge-1/2/0.0
user@R1# set interface lo0.0
```

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

```
user@R1# show security
forwarding-options {
  family iso {
    mode packet-based;
  }
}

user@R1# show interfaces
ge-1/2/0 {
  unit 0 {
    description to-R2;
    family inet {
      address 10.0.0.1/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.1/32;
    }
    family iso {
      address 49.0002.0192.0168.0001.00;
    }
  }
}

user@R1# show protocols
isis {
  interface ge-1/2/0.0;
  interface lo0.0;
}
```

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

Verification

Confirm that the configuration is working properly.

- [Verifying IS-IS Interface Configuration on page 17](#)
- [Verifying IS-IS Interface Configuration in Detail on page 17](#)
- [Verifying IS-IS Adjacencies on page 18](#)
- [Verifying IS-IS Adjacencies in Detail on page 18](#)

Verifying IS-IS Interface Configuration

Purpose Verify the status of the IS-IS-enabled interfaces.

Action From operational mode, enter the **show isis interface brief** command.

```
user@R1> show isis interface brief
```

IS-IS interface database:

Interface	L	CirID	Level 1	DR	Level 2	DR	L1/L2	Metric
lo0.0	0	0x1	Passive		Passive			0/0
ge-1/2/0.0	3	0x1	R2.02		R2.02		10/10	

Meaning Verify that the output shows the intended configuration of the interfaces on which IS-IS is enabled.

Verifying IS-IS Interface Configuration in Detail

Purpose Verify the details of IS-IS-enabled interfaces.

Action From operational mode, enter the **show isis interface detail** command.

```
user@R1> show isis interface detail
```

IS-IS interface database:

lo0.0

Index: 75, State: 0x6, Circuit id: 0x1, Circuit type: 0

LSP interval: 100 ms, CSNP interval: disabled

Adjacency advertisement: Advertise

Level	Adjacencies	Priority	Metric	Hello (s)	Hold (s)	Designated Router
1	0	64	0	Passive		
2	0	64	0	Passive		

ge-1/2/0.0

Index: 77, State: 0x6, Circuit id: 0x1, Circuit type: 3

LSP interval: 100 ms, CSNP interval: 10 s

Adjacency advertisement: Advertise

Level	Adjacencies	Priority	Metric	Hello (s)	Hold (s)	Designated Router
1	1	64	10	9.000	27	R2.02 (not us)
2	1	64	10	9.000	27	R2.02 (not us)

Meaning Check the following output fields and verify that the output shows the intended configuration of IS-IS-enabled interfaces:

- Interface—Interface configured for IS-IS.
- State—Internal implementation information.
- Circuit id—Circuit identifier.

- Circuit type—Configured level of IS-IS:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2
- link-state PDU interval—Time between IS-IS information messages.
- L or Level—Type of adjacency:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2
- Adjacencies—Adjacencies established on the interface.
- Priority—Priority value established on the interface.
- Metric—Metric value for the interface.
- Hello(s)—Intervals between hello PDUs.
- Hold(s)—Hold time on the interface.

Verifying IS-IS Adjacencies

Purpose Display brief information about IS-IS neighbors.

Action From operational mode, enter the **show isis adjacency brief** command.

```
user@R1> show isis adjacency brief
Interface      System      L State      Hold (secs) SNPA
ge-1/2/0.0     R2          1 Up          6 0:5:85:8f:c8:bd
ge-1/2/0.0     R2          2 Up          6 0:5:85:8f:c8:bd
```

Meaning Verify the adjacent routers in the IS-IS database.

Verifying IS-IS Adjacencies in Detail

Purpose Display extensive information about IS-IS neighbors.

Action From operational mode, enter the **show isis adjacency extensive** command.

```
user@R1> show isis adjacency extensive
R2
Interface: ge-1/2/0.0, Level: 1, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 00:40:28 ago
Circuit type: 3, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.0.2
Transition log:
When              State      Event              Down reason
Thu May 31 11:18:48 Up          Seenself
```

R2


```

Interface: ge-1/2/0.0, Level: 2, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 00:40:28 ago
Circuit type: 3, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.0.2
Transition log:
When                State      Event      Down reason
Thu May 31 11:18:48 Up         SeenseIf

```

Meaning Check the following fields and verify the adjacency information about IS-IS neighbors:

- Interface—Interface through which the neighbor is reachable.
- L or Level—Configured level of IS-IS:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2

An exclamation point before the level number indicates that the adjacency is missing an IP address.

- State—Status of the adjacency: **Up**, **Down**, **New**, **One-way**, **Initializing**, or **Rejected**.
- Event—Message that identifies the cause of a state.
- Down reason—Reason the adjacency is down.
- Restart capable—A neighbor is configured for graceful restart.
- Transition log—List of transitions including **When**, **State**, and **Reason**.

Related Documentation • [Example: Configuring Multi-Level IS-IS on page 19](#)

Example: Configuring Multi-Level IS-IS

- [Understanding IS-IS Areas on page 19](#)
- [Example: Configuring Multi-Level IS-IS on page 20](#)

Understanding IS-IS Areas

In IS-IS, a single AS can be divided into smaller groups called *areas*.

Link-state protocols cannot scale well if a large autonomous system (AS) consists of a single set of routing devices that all share a common database to compute the best paths through the AS. Because the shortest-path-first (SPF) algorithm works in an exponential fashion, the CPU demand can become too heavy when too many routing devices share their complete routing information with each other. To alleviate this issue, large ASs are divided into smaller parts called areas.

When ASs are split into areas, the disjointed areas must be connected to route traffic between the areas. Reachability information at the area borders must be injected into each other areas.

In IS-IS, routing between areas is organized hierarchically. 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, Level 1 systems route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs. No IS-IS area functions strictly as a backbone.

Level 1 routers share intra-area routing information, and Level 2 routers share interarea information about IP addresses available within each area. Uniquely, IS-IS routers can act as both Level 1 and Level 2 routers, sharing intra-area routes with other Level 1 routers and interarea routes with other Level 2 routers.

The propagation of link-state updates is determined by the level boundaries. All routers within a level maintain a complete link-state database of all other routers in the same level. Each router then uses the Dijkstra algorithm to determine the shortest path from the local router to other routers in the link-state database.

Example: Configuring Multi-Level IS-IS

This example shows how to configure a multi-level IS-IS topology.

- [Requirements on page 20](#)
- [Overview on page 20](#)
- [Configuration on page 21](#)
- [Verification on page 25](#)

Requirements

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

Overview


Like OSPF, the IS-IS protocol supports the partitioning of a routing domain into multiple areas with levels that control interarea flooding. The use of multiple levels improves protocol scalability, as Level 2 (backbone) link-state PDUs are normally not flooded into a Level 1 area.

An IS-IS Level 2 area is analogous to the OSPF backbone area (0), while a Level 1 area operates much like an OSPF totally stubby area, in that a default route is normally used to reach both inter-level and AS external routes.

Unlike OSPF, IS-IS area boundaries occur between routers, such that a given routing device is always wholly contained within a particular area. Level 1 adjacencies can be formed between routers that share a common area number, while a Level 2 adjacency can be formed between routers that might or might not share an area number.

[Figure 3 on page 21](#) shows the topology used in this example.

Figure 3: IS-IS Multi-Level Topology



"CLI Quick Configuration" on page 21 shows the configuration for all of the devices in Figure 3 on page 21. The section "Step-by-Step Procedure" on page 23 describes the steps on Device R5.

This example has the following characteristics:

- Device R5 functions as a Level 1/Level 2 router to interconnect the Level 2 backbone area 49.001 and the Level 1 area 49.002 containing Device R6 and Device R7.
- The system ID is based on the devices' IPv4 lo0 addresses.
- Loss of any individual interface does not totally disrupt the IS-IS operation.
- The IPv4 lo0 addresses of all routers are reachable through IS-IS.
- The link between Device R3 and Device S1 appears in area 49.001 as an intra-area route. No IS-IS adjacencies can be established on this interface. This is accomplished by configuring the **passive** statement on Device R3's interface to Device S1.
- The loopback addresses of Level 2 devices do not appear in a Level 1 area.
- There is only one adjacency for each device pairing.

Configuration

CLI Quick Configuration

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

Device R3

```
set interfaces fe-1/2/0 unit 0 description to-R4
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.21/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.41/30
```

```
set interfaces fe-1/2/2 unit 0 description to-S1
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set interfaces lo0 unit 0 family iso address 49.001.0192.0168.0003.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0 level 1 disable
set protocols isis interface fe-1/2/2.0 passive
```

Device R4

```
set interfaces fe-1/2/0 unit 0 description to-R3
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.18/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.25/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.001.0192.0168.0004.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0 level 1 disable
```

Device R5

```
set interfaces fe-1/2/0 unit 0 description to-R3
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.22/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R4
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.26/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 description to-R6
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.29/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/3 unit 0 description to-R7
set interfaces fe-1/2/3 unit 0 family inet address 10.0.0.38/30
set interfaces fe-1/2/3 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.5/32
set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0005.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface fe-1/2/2.0 level 2 disable
set protocols isis interface fe-1/2/3.0 level 2 disable
set protocols isis interface lo0.0 level 1 disable
```

Device R6

```
set interfaces fe-1/2/0 unit 0 description to-R5
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.30/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R7
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.33/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.6/32
set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0006.00
set protocols isis interface fe-1/2/0.0 level 2 disable
set protocols isis interface fe-1/2/1.0 level 2 disable
set protocols isis interface lo0.0 level 2 disable
```

Device R7

```
set interfaces fe-1/2/0 unit 0 description to-R6
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.34/30
set interfaces fe-1/2/0 unit 0 family iso
```

```

set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.37/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.7/32
set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0007.00
set protocols isis interface fe-1/2/0.0 level 2 disable
set protocols isis interface fe-1/2/1.0 level 2 disable
set protocols isis interface lo0.0 level 2 disable

```

Device S1 `set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.42/30`
`set interfaces fe-1/2/0 unit 0 description to-R3`

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

To configure multi-level IS-IS:

1. Configure the network interfaces.

Enable IS-IS on the interfaces by including the ISO address family on each interface.

```

[edit interfaces]
user@R5# set fe-1/2/0 unit 0 description to-R3
user@R5# set fe-1/2/0 unit 0 family inet address 10.0.0.22/30
user@R5# set fe-1/2/0 unit 0 family iso
user@R5# set fe-1/2/1 unit 0 description to-R4
user@R5# set fe-1/2/1 unit 0 family inet address 10.0.0.26/30
user@R5# set fe-1/2/1 unit 0 family iso
user@R5# set fe-1/2/2 unit 0 description to-R6
user@R5# set fe-1/2/2 unit 0 family inet address 10.0.0.29/30
user@R5# set fe-1/2/2 unit 0 family iso
user@R5# set fe-1/2/3 unit 0 description to-R7
user@R5# set fe-1/2/3 unit 0 family inet address 10.0.0.38/30
user@R5# set fe-1/2/3 unit 0 family iso

```

2. Configure two loopback interface addresses.

One address is for IPv4.

The other is for the IS-IS area 49.002 so that Device R5 can form adjacencies with the other Level 1 devices in area 49.002. Even though Device R5's NET identifies itself as belonging to the Level 1 area 49.002, its loopback interface is not configured as a Level 1 interface. Doing so would cause the route to Device R5's loopback to be injected into the Level 1 area.

```

[edit interfaces lo0 unit 0]
user@R5# set family inet address 192.168.0.5/32
user@R5# set family iso address 49.002.0192.0168.0005.00

```

3. Specify the IS-IS level on a per-interface basis.

Device R5 becomes adjacent to the other routing devices on the same level on each link.

By default, IS-IS is enabled for IS-IS areas on all interfaces on which the ISO protocol family is enabled (at the `[edit interfaces interface-name unit logical-unit-number]`

hierarchy level). To disable IS-IS at any particular level on an interface, include the **disable** statement.

Device R5's loopback interface is configured to run Level 2 only. If Level 1 operation were enabled on lo0.0, Device R5 would include its loopback address in its Level 1 link-state PDU, which is incorrect for this example in which the loopback addresses of Level 2 devices must not appear in a Level 1 area.

Unlike OSPF, you must explicitly list the router's lo0 interface at the **[edit protocols isis]** hierarchy level, because this interface is the source of the router's NET, and therefore must be configured as an IS-IS interface. In IS-IS, the lo0 interface operates in the passive mode by default, which is ideal because adjacency formation can never occur on a virtual interface.

```
[edit protocols isis]
user@R5# set interface fe-1/2/0.0 level 1 disable
user@R5# set interface fe-1/2/1.0 level 1 disable
user@R5# set interface fe-1/2/0.0 level 2 disable
user@R5# set interface fe-1/2/3.0 level 2 disable
user@R5# set interface lo0.0 level 1 disable
```

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

```
user@R5# show interfaces
fe-1/2/0 {
  unit 0 {
    description to-R3;
    family inet {
      address 10.0.0.22/30;
    }
    family iso;
  }
}
fe-1/2/1 {
  unit 0 {
    description to-R4;
    family inet {
      address 10.0.0.26/30;
    }
    family iso;
  }
}
fe-1/2/2 {
  unit 0 {
    description to-R6;
    family inet {
      address 10.0.0.29/30;
    }
    family iso;
  }
}
fe-1/2/3 {
  unit 0 {
```

```

description to-R7;
family inet {
    address 10.0.0.38/30;
}
family iso;
}
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.5/32;
        }
        family iso {
            address 49.002.0192.0168.0005.00;
        }
    }
}
}

user@R5# show protocols
isis {
    interface fe-1/2/0.0 {
        level 1 disable;
    }
    interface fe-1/2/1.0 {
        level 1 disable;
    }
    interface fe-1/2/0.0 {
        level 2 disable;
    }
    interface fe-1/2/3.0 {
        level 2 disable;
    }
    interface lo0.0 {
        level 1 disable;
    }
}

```

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

Verification

Confirm that the configuration is working properly.

- [Checking Interface-to-Area Associations on page 25](#)
- [Verifying IS-IS Adjacencies on page 26](#)
- [Examining the IS-IS Database on page 27](#)

Checking Interface-to-Area Associations

Purpose Make sure that the interface-to-area associations are configured as expected.

Action From operational mode, enter the **show isis interface** command.

```

user@R5> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR          Level 2 DR          L1/L2 Metric

```

10.0.0	0	0x1 Disabled	Passive	0/0
fe-1/2/0.0	2	0x3 Disabled	R5.03	10/10
fe-1/2/1.0	2	0x2 Disabled	R5.02	10/10
fe-1/2/0.0	1	0x1 R6.02	Disabled	10/10
fe-1/2/3.0	1	0x4 R5.04	Disabled	10/10

Meaning The output shows that Device R5's interfaces have been correctly configured with the ISO family, and that the interfaces have been placed into the correct levels.

You can also see that Device R5 has elected itself as the designated intermediate system (DIS) on its broadcast-capable IS-IS interfaces.

Verifying IS-IS Adjacencies

Purpose Verify that the expected adjacencies have formed between Device R5 and its IS-IS neighbors.

Action From operational mode, enter the **show isis adjacency detail** command.

```
user@R5> show isis adjacency detail
```

```
R3
```

```
Interface: fe-1/2/0.0, Level: 2, State: Up, Expires in 25 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:31 ago
Circuit type: 2, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.03, IP addresses: 10.0.0.21
```

```
R4
```

```
Interface: fe-1/2/1.0, Level: 2, State: Up, Expires in 24 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:36 ago
Circuit type: 2, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.02, IP addresses: 10.0.0.25
```

```
R6
```

```
Interface: fe-1/2/0.0, Level: 1, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:20:24 ago
Circuit type: 1, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R6.02, IP addresses: 10.0.0.30
```

```
R7
```

```
Interface: fe-1/2/3.0, Level: 1, State: Up, Expires in 21 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:29 ago
Circuit type: 1, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.04, IP addresses: 10.0.0.37
```

Meaning These results confirm that Device R5 has two Level 2 adjacencies and two Level 1 adjacencies.

Examining the IS-IS Database

Purpose Because Device R5 is a L1/L2 attached router, examine the Level 1 link-state database associated with area 49.002 to confirm that loopback addresses from backbone routers are not being advertised into the Level 1 area.

Action From operational mode, enter the **show isis database detail** command.

```
user@R5> show isis database detail
```

```
IS-IS level 1 link-state database:
```

```
R5.00-00 Sequence: 0x19, Checksum: 0x7488, Lifetime: 727 secs
  IS neighbor: R5.04                      Metric:      10
  IS neighbor: R6.02                      Metric:      10
  IP prefix: 10.0.0.28/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.36/30                  Metric:      10 Internal Up
```

```
R5.04-00 Sequence: 0x14, Checksum: 0x2668, Lifetime: 821 secs
  IS neighbor: R5.00                      Metric:       0
  IS neighbor: R7.00                      Metric:       0
```

```
R6.00-00 Sequence: 0x17, Checksum: 0xa65, Lifetime: 774 secs
  IS neighbor: R6.02                      Metric:      10
  IS neighbor: R7.02                      Metric:      10
  IP prefix: 10.0.0.28/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.32/30                  Metric:      10 Internal Up
  IP prefix: 192.168.0.6/32                Metric:       0 Internal Up
```

```
R6.02-00 Sequence: 0x13, Checksum: 0xd1c0, Lifetime: 908 secs
  IS neighbor: R5.00                      Metric:       0
  IS neighbor: R6.00                      Metric:       0
```

```
R7.00-00 Sequence: 0x17, Checksum: 0xe39, Lifetime: 775 secs
  IS neighbor: R5.04                      Metric:      10
  IS neighbor: R7.02                      Metric:      10
  IP prefix: 10.0.0.32/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.36/30                  Metric:      10 Internal Up
  IP prefix: 192.168.0.7/32                Metric:       0 Internal Up
```

```
R7.02-00 Sequence: 0x13, Checksum: 0x404d, Lifetime: 966 secs
  IS neighbor: R6.00                      Metric:       0
  IS neighbor: R7.00                      Metric:       0
```

```
IS-IS level 2 link-state database:
```

```
R3.00-00 Sequence: 0x17, Checksum: 0x5f84, Lifetime: 1085 secs
  IS neighbor: R4.02                      Metric:      10
  IS neighbor: R5.03                      Metric:      10
  IP prefix: 10.0.0.16/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.20/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.40/30                  Metric:      10 Internal Up
  IP prefix: 192.168.0.3/32                Metric:       0 Internal Up
```

```
R4.00-00 Sequence: 0x17, Checksum: 0xab3a, Lifetime: 949 secs
  IS neighbor: R4.02                      Metric:      10
  IS neighbor: R5.02                      Metric:      10
  IP prefix: 10.0.0.16/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.24/30                  Metric:      10 Internal Up
  IP prefix: 192.168.0.4/32                Metric:       0 Internal Up
```

```
R4.02-00 Sequence: 0x14, Checksum: 0xf2a8, Lifetime: 1022 secs
IS neighbor: R3.00                      Metric:      0
IS neighbor: R4.00                      Metric:      0

R5.00-00 Sequence: 0x1f, Checksum: 0x20d7, Lifetime: 821 secs
IS neighbor: R5.02                      Metric:      10
IS neighbor: R5.03                      Metric:      10
IP prefix: 10.0.0.20/30                 Metric:      10 Internal Up
IP prefix: 10.0.0.24/30                 Metric:      10 Internal Up
IP prefix: 10.0.0.28/30                 Metric:      10 Internal Up
IP prefix: 10.0.0.32/30                 Metric:      20 Internal Up
IP prefix: 10.0.0.36/30                 Metric:      10 Internal Up
IP prefix: 192.168.0.5/32               Metric:       0 Internal Up
IP prefix: 192.168.0.6/32               Metric:      10 Internal Up
IP prefix: 192.168.0.7/32               Metric:      10 Internal Up

R5.02-00 Sequence: 0x14, Checksum: 0x6135, Lifetime: 977 secs
IS neighbor: R4.00                      Metric:       0
IS neighbor: R5.00                      Metric:       0

R5.03-00 Sequence: 0x14, Checksum: 0x1483, Lifetime: 1091 secs
IS neighbor: R3.00                      Metric:       0
IS neighbor: R5.00                      Metric:       0
```

Meaning This display indicates that Device R5's loopback interface is correctly configured to run Level 2 only. Had Level 1 operation been enabled on lo0.0, Device R5 would have then included its loopback address in its Level 1 link-state PDU.

You can also see that Device R5 has Level 2 link-state PDUs, received from its adjacent neighbors.

Like an OSPF totally stubby area, no backbone (Level 2) or external prefixes are leaked into a Level 1 area, by default. Level 1 prefixes are leaked up into the IS-IS backbone, however, as can be seen in Device R5's Level 2 link-state PDU.

Related Documentation

- [Example: Configuring IS-IS Designated Routers on page 39](#)

Example: Configuring IS-IS on Logical Systems Within the Same Router

- [Introduction to Logical Systems on page 28](#)
- [Example: Configuring IS-IS on Logical Systems Within the Same Router on page 30](#)

Introduction to Logical Systems

For many years, engineers have combined power supplies, routing hardware and software, forwarding hardware and software, and physical interfaces into a networking device known as a router. Networking vendors have created large routers and small routers, but all routers have been placed into service as individual devices. As a result, the router has been considered a single physical device for most of its history.

The concept of logical systems breaks with this tradition. With the Junos[®] operating system (Junos OS), you can partition a single router into multiple logical devices that perform independent routing tasks. Because logical systems perform a subset of the

tasks once handled by the main router, logical systems offer an effective way to maximize the use of a single routing or switching platform.




NOTE: Beginning with Junos OS Release 9.3, the logical router feature has been renamed logical system.

All configuration statements, operational commands, show command output, error messages, log messages, and SNMP MIB objects that contain the string `logical-router` have been changed to `logical-system`.

Traditionally, service provider network design requires multiple layers of switches and routers. These devices transport packet traffic between customers. As seen on the left side of [Figure 4 on page 29](#), access devices are connected to edge devices, which are in turn connected to core devices.

However, this complexity can lead to challenges in maintenance, configuration, and operation. To reduce such complexity, Juniper Networks supports logical systems. Logical systems perform a subset of the actions of the main router and have their own unique routing tables, interfaces, policies, and routing instances. As shown on the right side of [Figure 4 on page 29](#), a set of logical systems within a single router can handle the functions previously performed by several small routers.

Figure 4: Logical Systems Concepts



[Figure 5 on page 30](#) shows the Junos OS architecture without logical systems configured. [Figure 6 on page 30](#) shows the Junos OS architecture when logical systems are configured. Note that each logical system runs its own routing protocol process (`rpd`).

Figure 5: Junos OS Without Logical Systems

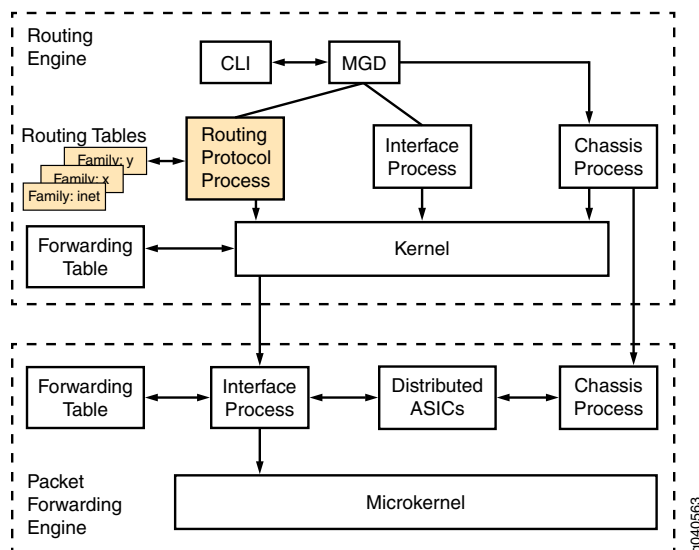
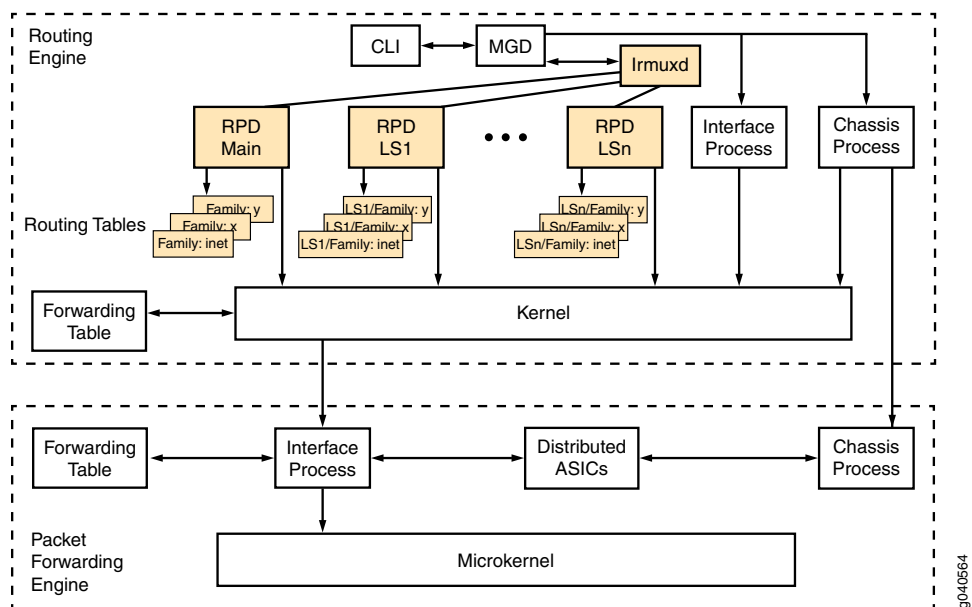


Figure 6: Junos OS With Logical Systems



Example: Configuring IS-IS on Logical Systems Within the Same Router

This example shows how to configure an IS-IS network by using multiple logical systems that are running on a single physical router. The logical systems are connected by logical tunnel interfaces.

- [Requirements on page 31](#)
- [Overview on page 31](#)
- [Configuration on page 31](#)
- [Verification on page 36](#)

Requirements


You must connect the logical systems by using logical tunnel (lt) interfaces. See *Example: Connecting Logical Systems Within the Same Router Using Logical Tunnel Interfaces*.

Overview

This example shows an IS-IS configuration with three logical systems running on one physical router. Each logical system has its own routing table. The configuration enables the protocol on all logical tunnel interfaces that participate in the IS-IS domain.

Figure 7 on page 31 shows the sample network.

Figure 7: IS-IS on Logical Systems



Configuration

CLI Quick Configuration

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

```
set logical-systems LS1 interfaces lt-0/1/0 unit 2 description LS1->LS2
set logical-systems LS1 interfaces lt-0/1/0 unit 2 encapsulation ethernet
set logical-systems LS1 interfaces lt-0/1/0 unit 2 peer-unit 1
set logical-systems LS1 interfaces lt-0/1/0 unit 2 family inet address 10.0.0.1/30
set logical-systems LS1 interfaces lt-0/1/0 unit 2 family iso
set logical-systems LS1 interfaces lt-0/1/0 unit 0 description LS1->LS3
set logical-systems LS1 interfaces lt-0/1/0 unit 0 encapsulation ethernet
set logical-systems LS1 interfaces lt-0/1/0 unit 0 peer-unit 5
set logical-systems LS1 interfaces lt-0/1/0 unit 0 family inet address 10.0.1.2/30
set logical-systems LS1 interfaces lt-0/1/0 unit 0 family iso
set logical-systems LS1 interfaces lo0 unit 1 family iso address 49.0001.1720.1600.1001.00
```

```

set logical-systems LS1 protocols isis interface lt-0/1/0.0
set logical-systems LS1 protocols isis interface lt-0/1/0.2
set logical-systems LS1 protocols isis interface lo0.1 passive
set logical-systems LS2 interfaces lt-0/1/0 unit 1 description LS2->LS1
set logical-systems LS2 interfaces lt-0/1/0 unit 1 encapsulation ethernet
set logical-systems LS2 interfaces lt-0/1/0 unit 1 peer-unit 2
set logical-systems LS2 interfaces lt-0/1/0 unit 1 family inet address 10.0.0.2/30
set logical-systems LS2 interfaces lt-0/1/0 unit 1 family iso
set logical-systems LS2 interfaces lt-0/1/0 unit 4 description LS2->LS3
set logical-systems LS2 interfaces lt-0/1/0 unit 4 encapsulation ethernet
set logical-systems LS2 interfaces lt-0/1/0 unit 4 peer-unit 3
set logical-systems LS2 interfaces lt-0/1/0 unit 4 family inet address 10.0.2.2/30
set logical-systems LS2 interfaces lt-0/1/0 unit 4 family iso
set logical-systems LS2 interfaces lo0 unit 2 family iso address
  49.0001.1720.1600.2002.00
set logical-systems LS2 protocols isis interface lt-0/1/0.1
set logical-systems LS2 protocols isis interface lt-0/1/0.4
set logical-systems LS2 protocols isis interface lo0.2 passive
set logical-systems LS3 interfaces lt-0/1/0 unit 3 description LS3->LS2
set logical-systems LS3 interfaces lt-0/1/0 unit 3 encapsulation ethernet
set logical-systems LS3 interfaces lt-0/1/0 unit 3 peer-unit 4
set logical-systems LS3 interfaces lt-0/1/0 unit 3 family inet address 10.0.2.1/30
set logical-systems LS3 interfaces lt-0/1/0 unit 3 family iso
set logical-systems LS3 interfaces lt-0/1/0 unit 5 description LS3->LS1
set logical-systems LS3 interfaces lt-0/1/0 unit 5 encapsulation ethernet
set logical-systems LS3 interfaces lt-0/1/0 unit 5 peer-unit 0
set logical-systems LS3 interfaces lt-0/1/0 unit 5 family inet address 10.0.1.1/30
set logical-systems LS3 interfaces lt-0/1/0 unit 5 family iso
set logical-systems LS3 interfaces lo0 unit 3 family iso address 49.0001.1234.1600.2231.00
set logical-systems LS3 protocols isis interface lt-0/1/0.5
set logical-systems LS3 protocols isis interface lt-0/1/0.3
set logical-systems LS3 protocols isis interface lo0.3 passive

```

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

To configure IS-IS on logical systems:

1. Configure the logical tunnel interface on Logical System LS1 connecting to Logical System LS2.

```

[edit logical-systems LS1]
user@host# set interfaces lt-0/1/0 unit 2 description LS1->LS2
user@host# set interfaces lt-0/1/0 unit 2 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 2 peer-unit 1
user@host# set interfaces lt-0/1/0 unit 2 family inet address 10.0.0.1/30
user@host# set interfaces lt-0/1/0 unit 2 family iso

```

2. Configure the logical tunnel interface on Logical System LS1 connecting to Logical System LS3.

```

[edit logical-systems LS1]
user@host# set interfaces lt-0/1/0 unit 0 description LS1->LS3
user@host# set interfaces lt-0/1/0 unit 0 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 0 peer-unit 5

```

- ```

user@host# set interfaces lt-0/1/0 unit 0 family inet address 10.0.1.2/30
user@host# set interfaces lt-0/1/0 unit 0 family iso

```
3. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS1.
 

```

[edit logical-systems LS2]
user@host# set interfaces lt-0/1/0 unit 1 description LS2->LS1
user@host# set interfaces lt-0/1/0 unit 1 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 1 peer-unit 2
user@host# set interfaces lt-0/1/0 unit 1 family inet address 10.0.0.2/30
user@host# set interfaces lt-0/1/0 unit 1 family iso

```
  4. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS3.
 

```

[edit logical-systems LS2]
user@host# set interfaces lt-0/1/0 unit 4 description LS2->LS3
user@host# set interfaces lt-0/1/0 unit 4 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 4 peer-unit 3
user@host# set interfaces lt-0/1/0 unit 4 family inet address 10.0.2.2/30
user@host# set interfaces lt-0/1/0 unit 4 family iso

```
  5. Configure the logical tunnel interface on Logical System LS3 connecting to Logical System LS2.
 

```

[edit logical-systems LS3]
user@host# set interfaces lt-0/1/0 unit 3 description LS3->LS2
user@host# set interfaces lt-0/1/0 unit 3 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 3 peer-unit 4
user@host# set interfaces lt-0/1/0 unit 3 family inet address 10.0.2.1/30
user@host# set interfaces lt-0/1/0 unit 3 family iso

```
  6. Configure the logical tunnel interface on Logical System LS3 connecting to Logical System LS1.
 

```

[edit logical-systems LS3]
user@host# set interfaces lt-0/1/0 unit 5 description LS3->LS1
user@host# set interfaces lt-0/1/0 unit 5 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 5 peer-unit 0
user@host# set interfaces lt-0/1/0 unit 5 family inet address 10.0.1.1/30
user@host# set interfaces lt-0/1/0 unit 5 family iso

```
  7. Configure the ISO address on the loopback interface for the three logical systems.
 

```

[edit logical-systems LS1]
user@host# set interfaces lo0 unit 1 family iso address 49.0001.1720.1600.1001.00
user@host# set protocols isis interface lo0.1 passive

[edit logical-systems LS2]
user@host# set interfaces lo0 unit 2 family iso address 49.0001.1720.1600.2002.00
user@host# set protocols isis interface lo0.2 passive

[edit logical-systems LS3]
user@host# set interfaces lo0 unit 3 family iso address 49.0001.1234.1600.2231.00
user@host# set protocols isis interface lo0.3 passive

```
  8. Configure IS-IS on all the interfaces.
 

```

[edit logical-systems LS1 protocols isis]
user@host# set interface lt-0/1/0.0

```

```
user@host# set interface lt-0/1/0.2
[edit logical-systems LS2 protocols isis]
user@host# set interface lt-0/1/0.1
user@host# set interface lt-0/1/0.4
[edit logical-systems LS3 protocols isis]
user@host# set interface lt-0/1/0.5
user@host# set interface lt-0/1/0.3
```

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

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

### Results

From configuration mode, confirm your configuration by issuing the **show logical-systems** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show logical-systems
LS1 {
 interfaces {
 lt-0/1/0 {
 unit 0 {
 description LS1->LS3;
 encapsulation ethernet;
 peer-unit 5;
 family inet {
 address 10.0.1.2/30;
 }
 family iso;
 }
 unit 2 {
 description LS1->LS2;
 encapsulation ethernet;
 peer-unit 1;
 family inet {
 address 10.0.0.1/30;
 }
 family iso;
 }
 }
 lo0 {
 unit 1 {
 family iso {
 address 49.0001.1720.1600.1001.00;
 }
 }
 }
 }
 protocols {
 isis {
 interface lt-0/1/0.0;
 interface lt-0/1/0.2;
 interface lo0.1 {
 passive;
 }
 }
 }
}
```



```

}
LS2 {
 interfaces {
 lt-0/1/0 {
 unit 1 {
 description LS2->LS1;
 encapsulation ethernet;
 peer-unit 2;
 family inet {
 address 10.0.0.2/30;
 }
 family iso;
 }
 unit 4 {
 description LS2->LS3;
 encapsulation ethernet;
 peer-unit 3;
 family inet {
 address 10.0.2.2/30;
 }
 family iso;
 }
 }
 lo0 {
 unit 2 {
 family iso {
 address 49.0001.1720.1600.2002.00;
 }
 }
 }
 }
 protocols {
 isis {
 interface lt-0/1/0.1;
 interface lt-0/1/0.4;
 interface lo0.2 {
 passive;
 }
 }
 }
}
LS3 {
 interfaces {
 lt-0/1/0 {
 unit 3 {
 description LS3->LS2;
 encapsulation ethernet;
 peer-unit 4;
 family inet {
 address 10.0.2.1/30;
 }
 family iso;
 }
 unit 5 {
 description LS3->LS1;
 encapsulation ethernet;
 peer-unit 0;
 family inet {
 address 10.0.1.1/30;
 }
 }
 }
 }
}

```

```

 family iso;
 }
}
lo0 {
 unit 3 {
 family iso {
 address 49.0001.1234.1600.2231.00;
 }
 }
}
}
protocols {
 isis {
 interface lt-0/1/0.3;
 interface lt-0/1/0.5;
 interface lo0.3 {
 passive;
 }
 }
}
}

```

### Verification

Confirm that the configuration is working properly.

- [Verifying That the Logical Systems Are Up on page 36](#)
- [Verifying Connectivity Between the Logical Systems on page 36](#)

#### *Verifying That the Logical Systems Are Up*

**Purpose** Make sure that the interfaces are properly configured.

|               |                                  |       |      |       |             |        |
|---------------|----------------------------------|-------|------|-------|-------------|--------|
| <b>Action</b> | user@host> show interfaces terse |       |      |       |             |        |
|               | Interface                        | Admin | Link | Proto | Local       | Remote |
|               | ...                              |       |      |       |             |        |
|               | lt-0/1/0                         | up    | up   |       |             |        |
|               | lt-0/1/0.0                       | up    | up   | inet  | 10.0.1.2/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | lt-0/1/0.1                       | up    | up   | inet  | 10.0.0.2/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | lt-0/1/0.2                       | up    | up   | inet  | 10.0.0.1/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | lt-0/1/0.3                       | up    | up   | inet  | 10.0.2.1/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | lt-0/1/0.4                       | up    | up   | inet  | 10.0.2.2/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | lt-0/1/0.5                       | up    | up   | inet  | 10.0.1.1/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | ...                              |       |      |       |             |        |

#### *Verifying Connectivity Between the Logical Systems*

**Purpose** Make sure that the IS-IS adjacencies are established by checking the logical system routing entries and by pinging the logical systems.

```

Action user@host> show route logical-system LS1
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30 *[Direct/0] 3w0d 01:37:52
 > via lt-0/1/0.2
10.0.0.1/32 *[Local/0] 3w0d 01:37:52
 Local via lt-0/1/0.2
10.0.1.0/30 *[Direct/0] 3w0d 01:37:52
 > via lt-0/1/0.0
10.0.1.2/32 *[Local/0] 3w0d 01:37:52
 Local via lt-0/1/0.0
10.0.2.0/30 *[IS-IS/15] 3w0d 01:37:13, metric 20
 > to 10.0.1.1 via lt-0/1/0.0
 to 10.0.0.2 via lt-0/1/0.2

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.1001/72
 *[Direct/0] 3w0d 01:37:52
 > via lo0.1

user@host> show route logical-system LS2
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30 *[Direct/0] 3w0d 01:38:01
 > via lt-0/1/0.1
10.0.0.2/32 *[Local/0] 3w0d 01:38:01
 Local via lt-0/1/0.1
10.0.1.0/30 *[IS-IS/15] 3w0d 01:37:01, metric 20
 to 10.0.0.1 via lt-0/1/0.1
 > to 10.0.2.1 via lt-0/1/0.4
10.0.2.0/30 *[Direct/0] 3w0d 01:38:01
 > via lt-0/1/0.4
10.0.2.2/32 *[Local/0] 3w0d 01:38:01
 Local via lt-0/1/0.4

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.2002/72
 *[Direct/0] 3w0d 01:38:01
 > via lo0.2

user@host> show route logical-system LS3
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30 *[IS-IS/15] 3w0d 01:37:10, metric 20
 to 10.0.2.2 via lt-0/1/0.3
 > to 10.0.1.2 via lt-0/1/0.5
10.0.1.0/30 *[Direct/0] 3w0d 01:38:10
 > via lt-0/1/0.5
10.0.1.1/32 *[Local/0] 3w0d 01:38:11
 Local via lt-0/1/0.5
10.0.2.0/30 *[Direct/0] 3w0d 01:38:11
 > via lt-0/1/0.3

```

```

10.0.2.1/32 *[Local/0] 3w0d 01:38:11
 Local via lt-0/1/0.3

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1234.1600.2231/72
 *[Direct/0] 3w0d 01:38:11
 > via lo0.3

```

#### From LS1, Ping LS3

```

user@host> set cli logical-system LS1

user@host:LS1> ping 10.0.2.1
PING 10.0.2.1 (10.0.2.1): 56 data bytes
64 bytes from 10.0.2.1: icmp_seq=0 ttl=63 time=1.264 ms
64 bytes from 10.0.2.1: icmp_seq=1 ttl=63 time=1.189 ms
64 bytes from 10.0.2.1: icmp_seq=2 ttl=63 time=1.165 ms
^C
--- 10.0.2.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.165/1.206/1.264/0.042 ms

```

#### From LS3, Ping LS1

```

user@host> set cli logical-system LS3

user@host:LS3> ping 10.0.0.1
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=63 time=1.254 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=63 time=1.210 ms
^C
--- 10.0.0.1 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.210/1.232/1.254/0.022 ms

```

#### From LS1, Ping LS2

```

user@host> set cli logical-system LS1

user@host:LS1> ping 10.0.2.2
PING 10.0.2.2 (10.0.2.2): 56 data bytes
64 bytes from 10.0.2.2: icmp_seq=0 ttl=64 time=1.240 ms
64 bytes from 10.0.2.2: icmp_seq=1 ttl=64 time=1.204 ms
64 bytes from 10.0.2.2: icmp_seq=2 ttl=64 time=1.217 ms
^C
--- 10.0.2.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.204/1.220/1.240/0.015 ms

```

#### From LS2, Ping LS1

```

user@host> set cli logical-system LS2

user@host:LS2> ping 10.0.1.2
PING 10.0.1.2 (10.0.1.2): 56 data bytes
64 bytes from 10.0.1.2: icmp_seq=0 ttl=64 time=1.308 ms
64 bytes from 10.0.1.2: icmp_seq=1 ttl=64 time=1.235 ms
^C
--- 10.0.1.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.235/1.272/1.308/0.037 ms

```

### From LS2, Ping LS3

```
user@host> set cli logical-system LS2

user@host:LS2> ping 10.0.1.1
PING 10.0.1.1 (10.0.1.1): 56 data bytes
64 bytes from 10.0.1.1: icmp_seq=0 ttl=64 time=1.253 ms
64 bytes from 10.0.1.1: icmp_seq=1 ttl=64 time=1.194 ms
64 bytes from 10.0.1.1: icmp_seq=2 ttl=64 time=1.212 ms
64 bytes from 10.0.1.1: icmp_seq=3 ttl=64 time=1.221 ms
64 bytes from 10.0.1.1: icmp_seq=4 ttl=64 time=1.195 ms
^C
--- 10.0.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.194/1.215/1.253/0.022 ms
```

### From LS3, Ping LS2

```
user@host> set cli logical-system LS3

user@host:LS3> ping 10.0.0.2
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=64 time=1.240 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=1.217 ms
^C
--- 10.0.0.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.217/1.228/1.240/0.012 ms
```

**Related Documentation**

- [Example: Configuring IS-IS on page 13](#)

## Example: Configuring IS-IS Designated Routers

- [Understanding IS-IS Designated Routers on page 39](#)
- [Example: Configuring Designated Router Election Priority for IS-IS on page 40](#)

### Understanding IS-IS Designated Routers

A router advertises its priority to become a designated router in its hello packets. On all multiaccess networks (physical networks that support the attachment of more than two routers, such as Ethernet networks), IS-IS uses the advertised priorities to elect a designated router for the network. This router is responsible for sending network link-state advertisements, which describe all the routers attached to the network. These advertisements are flooded throughout a single area. The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.

A router's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127, which you configure on the IS-IS interface. The router with the highest priority becomes the designated router for the area (Level 1, Level 2, or both), also configured on the IS-IS interface. If routers in the network have the same priority, then the router with the highest MAC address is elected as the designated router. By default, routers have a priority value of 64.

## Example: Configuring Designated Router Election Priority for IS-IS

This example shows how to configure the designated router election priority for IS-IS.

Before you begin:

- Configure network interfaces. See the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Enable IS-IS on the interfaces. See [“Example: Configuring IS-IS” on page 14](#).

In this example, you configure the priority for logical interface ge-0/0/1.0 to be 100 and the level number to be 1. If this interface has the highest priority value, the router becomes the designated router for the Level 1 area.

To configure a designated router election priority for IS-IS:

```
[edit]
user@host# set protocols isis interface ge-0/0/1.0 level 1 priority 100
```

Related  
Documentation

- [Example: Configuring IS-IS on page 13](#)

---

## Configuring an ISO System Identifier for the Router

For IS-IS to operate on the router, you can optionally configure a system identifier (system ID). The system identifier is commonly the media access control (MAC) address or the IP address expressed in binary-coded decimal (BCD).

If you do not statically map the hostname, the mapping is generated dynamically, based on the system host-name. If you omit the **static-host-mapping hostname sysid** statement, the IS-IS system ID is dynamically generated from the host portion of the ISO address configured on the loopback interface (lo0) and is mapped to the **host-name** statement configured at the **[edit system]** hierarchy level. Run the [show isis hostname](#) command to view the mappings.

To configure an International Organization for Standardization (ISO) system ID, include the **sysid** statement at the **[edit system static-host-mapping hostname]** hierarchy level:

```
[edit system]
static-host-mapping {
 hostname {
 sysid system-identifier;
 }
}
```

**hostname** is the name specified by the **host-name** statement at the **[edit system]** hierarchy level.

**system-identifier** is the ISO system identifier. It is the 6-byte system ID portion of the IS-IS network service access point (NSAP). We recommend that you use the host's IP address represented in BCD format. For example, the IP address 192.168.1.77 is 1921.6800.1077 in BCD.

- Related Documentation**
- [Example: Configuring the Name of the Router, IP Address, and System ID on page 41](#)

## Example: Configuring the Name of the Router, IP Address, and System ID

The following example shows how to configure the router's name, map the name to an IP address and alias, and configure a system identifier:

```
[edit]
user@host# set system host-name router-sjl
[edit]
user@host# set system static-host-mapping router-sjl inet 192.168.1.77
[edit]
user@host# set system static-host-mapping router-sjl alias sjl
[edit]
user@host# set system static-host-mapping router-sjl sysid 1921.6800.1077
[edit]
user@host# show
system {
 host-name router-sjl;
 static-host-mapping {
 router-sjl {
 inet 192.168.1.77;
 alias sjl;
 sysid 1921.6800.1077;
 }
 }
}
```

- Related Documentation**
- [Configuring an ISO System Identifier for the Router on page 40](#)





# IS-IS Authentication and Checksums

- [Configuring IS-IS Authentication on page 43](#)
- [Configuring Authentication Without Network-Wide Deployment on page 45](#)
- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 45](#)
- [Example: Enabling Packet Checksums on IS-IS Interfaces on page 50](#)

## Configuring IS-IS Authentication

---

All IS-IS protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in the autonomous system (AS) routing. By default, IS-IS authentication is disabled on the routing device.

To configure IS-IS authentication, you must define an authentication password and specify the authentication type.

You can configure one of the following authentication methods:

- Simple authentication—Uses a text password that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet. Simple authentication is included for compatibility with existing IS-IS implementations. However, we recommend that you do *not* use this authentication method because it is insecure (the text can be “sniffed”).



**CAUTION:** A simple password that exceeds 254 characters is truncated.

- HMAC-MD5 authentication—Uses an iterated cryptographic hash function. The receiving routing device uses an authentication key (password) to verify the packet.

You can also configure more fine-grained interface-level authentication for hello packets.

To enable authentication and specify an authentication method, include the **authentication-type** statement, specifying the **simple** or **md5** authentication type:

**authentication-type** *authentication*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To configure a password, include the **authentication-key** statement. The authentication password for all routing devices in a domain must be the same.

**authentication-key** *key*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To configure hitless authentication key rollover, include the **authentication-key-chain (Protocols IS-IS)** statement.

The password can contain up to 255 characters. If you include spaces, enclose all characters in quotation marks (" ").

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 that are shared with a Junos OS implementation.

Authentication of hello packets, partial sequence number PDU (PSNP), and complete sequence number PDU (CSNP) can be suppressed to enable interoperability with the routing software of different vendors. Different vendors handle authentication in various ways, and suppressing authentication for different PDU types might be the simplest way to allow compatibility within the same network.

To configure IS-IS to generate authenticated packets, but not to check the authentication on received packets, include the **no-authentication-check** statement:

**no-authentication-check**;

To suppress authentication of IS-IS hello packets, include the **no-hello-authentication** statement:

**no-hello-authentication**;

To suppress authentication of PSNPs, include the **no-psnp-authentication** statement:

**no-psnp-authentication**;

To suppress authentication of CSNPs, include the **no-csnp-authentication** statement:

**no-csnp-authentication**;

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.



**NOTE:** The **authentication** and the **no-authentication** statements must be configured at the same hierarchy level. Configuring authentication at the [edit protocols isis interface *interface-name*] hierarchy level and configuring **no-authentication** at the [edit protocols isis] hierarchy level has no effect.

---

#### Related Documentation

- [Configuring Authentication Without Network-Wide Deployment on page 45](#)

## Configuring Authentication Without Network-Wide Deployment

To allow the use of authentication without requiring network-wide deployment, include the **loose-authentication-check** statement:

```
loose-authentication-check;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Related Documentation

- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 45](#)

## Example: Configuring Hitless Authentication Key Rollover for IS-IS

- [Understanding Hitless Authentication Key Rollover for IS-IS on page 45](#)
- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 46](#)

### Understanding Hitless Authentication Key Rollover for IS-IS

IS-IS protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in routing. By default, authentication is disabled. The authentication algorithm creates an encoded checksum that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet's checksum.

If you configure authentication for all peers, each peer in that group inherits the group's authentication.

You can update authentication keys without resetting any IS-IS neighbor sessions. This is referred to as *hitless authentication key rollover*.

Hitless authentication key rollover uses authentication keychains, which consist of the authentication keys that are being updated. The keychain includes multiple keys. Each key in the keychain has a unique start time. At the next key's start time, a rollover occurs from the current key to the next key, and the next key becomes the current key.

You can choose the algorithm through which authentication is established. You can configure MD5 or SHA-1 authentication. You associate a keychain and the authentication algorithm with an IS-IS neighboring session. Each key contains an identifier and a secret password.

The sending peer chooses the active key based on the system time and the start times of the keys in the keychain. The receiving peer determines the key with which it authenticates based on the incoming key identifier.

You can configure either RFC 5304-based encoding or RFC 5310-based encoding for the IS-IS protocol transmission encoding format.

## Example: Configuring Hitless Authentication Key Rollover for IS-IS

This example shows how to configure hitless authentication key rollover for IS-IS.

- [Requirements on page 46](#)
- [Overview on page 46](#)
- [Configuration on page 47](#)
- [Verification on page 50](#)

---

### Requirements

No special configuration beyond device initialization is required before configuring hitless authentication key rollover for IS-IS.

---

### Overview

Authentication guarantees that only trusted routers participate in routing updates. This keychain authentication method is referred to as hitless because the keys roll over from one to the next without resetting any peering sessions or interrupting the routing protocol. Junos OS supports both RFC 5304, *IS-IS Cryptographic Authentication* and RFC 5310, *IS-IS Generic Cryptographic Authentication*.

This example includes the following statements for configuring the keychain:

- **algorithm**—For each key in the keychain, you can specify an encryption algorithm. The algorithm can be SHA-1 or MD-5.
- **key**—A keychain can have multiple keys. Each key within a keychain must be identified by a unique integer value. The range of valid identifier values is from 0 through 63.
- **key-chain**—For each keychain, you must specify a name. This example defines two keychains: **base-key-global** and **base-key-inter**.
- **options**—For each key in the keychain, you can specify the encoding for the message authentication code: **isis-enhanced** or **basic**. The basic (RFC 5304) operation is enabled by default.

When you configure the **isis-enhanced** option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.

When you configure **basic** (or do not include the **options** statement in the key configuration), Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices.

Because this setting is for IS-IS only, the TCP and the BFD protocols ignore the encoding option configured in the key.

- **secret**—For each key in the keychain, you must set a secret password. This password can be entered in either encrypted or plain text format in the **secret** statement. It is always displayed in encrypted format.
- **start-time**—Each key must specify a start time in UTC format. Control gets passed from one key to the next. When a configured start time arrives (based on the routing device's clock), the key with that start time becomes active. Start times are specified in the local time zone for a routing device and must be unique within the key chain.


You can apply a keychain globally to all interfaces or more granularly to specific interfaces.

This example includes the following statements for applying the keychain to all interfaces or to particular interfaces:

- **authentication-key-chain**—Enables you to apply a keychain at the global IS-IS level for all Level 1 or all Level 2 interfaces.
- **hello-authentication-key-chain**—Enables you to apply a keychain at the individual IS-IS interface level. The interface configuration overrides the global configuration.

Figure 8 on page 47 shows the topology used in the example.

**Figure 8: Hitless Authentication Key Rollover for IS-IS**



This example shows the configuration for Router R0.

### Configuration

#### CLI Quick Configuration

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

```
set interfaces ge-0/0/0 unit 0 description "interface A"
set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.1/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family inet6 address fe80::200:f8ff:fe21:67cf/128
set interfaces ge-0/0/1 unit 0 description "interface B"
set interfaces ge-0/0/1 unit 0 family inet address 10.0.0.5/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family inet6 address 10FB::C:ABC:1FOC:44DA/128
set interfaces ge-0/0/2 unit 0 description "interface C"
set interfaces ge-0/0/2 unit 0 family inet address 10.0.0.9/30
set interfaces ge-0/0/2 unit 0 family iso
```

```

set interfaces ge-0/0/2 unit 0 family inet6 address ff06::c3/128
set security authentication-key-chains key-chain base-key-global key 63 secret
"9JfkqfTQnCpBDiCt"
set security authentication-key-chains key-chain base-key-global key 63 start-time
"2011-8-6.06:54:00-0700"
set security authentication-key-chains key-chain base-key-global key 63 algorithm
hmac-sha-1
set security authentication-key-chains key-chain base-key-global key 63 options
isis-enhanced
set security authentication-key-chains key-chain base-key-inter key 0 secret
"$9$8sgx7Vws4ZDkWLGD"
set security authentication-key-chains key-chain base-key-inter key 0 start-time
"2011-8-6.06:54:00-0700"
set security authentication-key-chains key-chain base-key-inter key 0 algorithm md5
set security authentication-key-chains key-chain base-key-inter key 0 options basic
set protocols isis level 1 authentication-key-chain base-key-global
set protocols isis interface ge-0/0/0.0 level 1 hello-authentication-key-chain
base-key-inter

```

**Step-by-Step Procedure** To configure hitless authentication key rollover for IS-IS:

1. Configure the Router R0 interfaces.

```

[edit interfaces ge-0/0/0 unit 0]
user@R0# set description "interface A"
user@R0# set family inet address 10.0.0.1/30
user@R0# set family iso
user@R0# set family inet6 address fe80::200:f8ff:fe21:67cf/128
[edit interfaces ge-0/0/1 unit 0]
user@R0# set interfaces ge-0/0/1 unit 0 description "interface B"
user@R0# set interfaces ge-0/0/1 unit 0 family inet address 10.0.0.5/30
user@R0# set interfaces ge-0/0/1 unit 0 family iso
user@R0# set interfaces ge-0/0/1 unit 0 family inet6 address
10FB::C:ABC:1FOC:44DA/128
[edit interfaces ge-0/0/2 unit 0]
user@R0# set description "interface C"
user@R0# set family inet address 10.0.0.9/30
user@R0# set interfaces ge-0/0/2 unit 0 family iso
user@R0# set interfaces ge-0/0/2 unit 0 family inet6 address ff06::c3/128

```

2. Configure one or more authentication keys.

```

[edit security authentication-key-chains key-chain base-key-global]
user@R0# set key 63 secret "9JfkqfTQnCpBDiCt"
user@R0# set key 63 start-time "2011-8-6.06:54:00-0700"
user@R0# set key 63 algorithm hmac-sha-1
user@R0# set key 63 options isis-enhanced
[edit security authentication-key-chains key-chain base-key-inter]
user@R0# set key 0 secret "$9$8sgx7Vws4ZDkWLGD"
user@R0# set key 0 start-time "2011-8-6.06:54:00-0700"
user@R0# set key 0 algorithm md5
user@R0# set key 0 options basic

```

3. Apply the base-key-global keychain to all Level 1 IS-IS interfaces on Router R0.

```

[edit protocols isis level 1]
user@R0# set authentication-key-chain base-key-global

```

4. Apply the base-key-inter keychain to the ge-0/0/0.0 interface on Router R0.

```
[edit protocols isis interface ge-0/0/0.0 level 1]
user@R0# set hello-authentication-key-chain base-key-inter
```

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

```
user@R0# commit
```

### Results

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

```
user@R0# show interfaces
ge-0/0/0 {
 unit 0 {
 description "interface A";
 family inet {
 address 10.0.0.1/30;
 }
 family iso;
 family inet6 {
 address fe80::200:f8ff:fe21:67cf/128;
 }
 }
}
ge-0/0/1 {
 unit 0 {
 description "interface B";
 family inet {
 address 10.0.0.5/30;
 }
 family iso;
 family inet6 {
 address 10fb::c:abc:1f0c:44da/128;
 }
 }
}
ge-0/0/2 {
 unit 0 {
 description "interface C";
 family inet {
 address 10.0.0.9/30;
 }
 family iso;
 family inet6 {
 address ff06::c3/128;
 }
 }
}

user@R0# show protocols
isis {
 level 1 authentication-key-chain base-key-global;
```

```
interface ge-0/0/0.0 {
 level 1 hello-authentication-key-chain base-key-inter;
}
}

user@R0# show security
authentication-key-chains {
 key-chain base-key-global {
 key 63 {
 secret "9jfkqfTQnCpBDiCt"; ## SECRET-DATA
 start-time "2011-8-6.06:54:00-0700";
 algorithm hmac-sha-1;
 options isis-enhanced;
 }
 }
 key-chain base-key-inter {
 key 0 {
 secret "$9$8sgx7Vws4ZDkWLGD"; ## SECRET-DATA
 start-time "2011-8-6.06:54:00-0700";
 algorithm md5;
 options basic;
 }
 }
}
```

---

### Verification

To verify the configuration, run the following commands:

- [show isis authentication](#)
- [show security keychain](#)

### Related Documentation

- [Example: Configuring Link and Node Protection for IS-IS Routes on page 141](#)

---

## Example: Enabling Packet Checksums on IS-IS Interfaces

- [Understanding Checksums on IS-IS Interfaces on page 50](#)
- [Example: Enabling Packet Checksums on IS-IS Interfaces on page 51](#)

### Understanding Checksums on IS-IS Interfaces

The checksum enables IS-IS to check at the receiver if the IS-IS protocol frames have become corrupted while traversing the network.

Sometimes corrupt IS-IS protocol frames can go undetected. If routing control traffic becomes corrupted, it is likely that user payload traffic might be corrupted, too. This can lead to unacceptable throughput. To prevent corrupt frames from going undetected, we recommend enabling checksumming on the IS-IS interfaces.

To review, IS-IS hello (IIH) PDUs establish adjacencies with other routing devices. A partial sequence number PDU (PSNP) is used by an IS-IS router to request link-state



PDU information from a neighboring router. The complete sequence number PDU (CSNP) lists all the link-state PDUs in the link-state database.

The original specification for IS-IS does not provide checksums for IIHs, CSNPs, and PSNPs.

RFC 3358, *Optional Checksums in Intermediate System to Intermediate System (ISIS)* introduced an optional type, length, and value (TLV) tuple that provides checksums for IIHs, PSNPs, and CSNPs.

Junos OS supports the checksumming TLV on a per-interface basis.

## Example: Enabling Packet Checksums on IS-IS Interfaces

This example shows how to enable packet checksums for IS-IS interfaces.

- [Requirements on page 51](#)
- [Overview on page 51](#)
- [Configuration on page 52](#)
- [Verification on page 52](#)

### Requirements

Before you begin, configure IS-IS on both routers. See “[Example: Configuring IS-IS](#)” on [page 14](#) for information about the sample IS-IS configuration.

### Overview


Junos OS supports IS-IS checksums as documented in RFC 3358, *Optional Checksums in Intermediate System to Intermediate System (ISIS)*.

IS-IS protocol data units (PDUs) include link-state PDUs, complete sequence number PDUs (CSNPs), partial sequence number PDUs (PSNPs), and IS-IS hello (IIH) packets. These PDUs can be corrupt due to faulty implementations of Layer 2 hardware or lack of checksums on a specific network technology. Corruption of length or type, length, and value (TLV) fields can lead to the generation of extensive numbers of empty link-state PDUs in the receiving node. Because authentication is not a replacement for a checksum mechanism, you might want to enable the optional checksum TLV on your IS-IS interfaces.

The checksum cannot be enabled with MD5 hello authentication on the same interface.

[Figure 9 on page 51](#) shows the topology used in this example.

**Figure 9: IS-IS Checksum Topology**



This example describes the steps on Device R1.

### Configuration

---

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

**Device R1**

```
set protocols isis traceoptions file isis
set protocols isis traceoptions flag all
set protocols isis interface fe-1/2/0.1 checksum
```

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

To configure IS-IS checksums:

1. Enable checksums.

```
[edit protocols isis interface fe-1/2/0.1]
user@R1# set checksum
```

2. (Optional) Enable tracing for tracking checksum operations.

```
[edit protocols isis traceoptions]
user@R1# set file isis
user@R1# set flag all
```

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

```
user@R1# show protocols
isis {
 traceoptions {
 file isis;
 flag all;
 }
 interface fe-1/2/0.1 {
 checksum;
 }
}
```

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

### Verification

---

Confirm that the configuration is working properly.

#### Verifying Checksums

**Purpose** Verify that checksums are performed.

**Action** From operational mode, enter the `show log isis | match checksum` command.

```
user@R1> show log isis | match checksum
```

```
May 31 16:47:39.513267 sequence 0x49 checksum 0x8e64
May 31 16:47:39.513394 sequence 0x4e checksum 0x34b3
May 31 16:47:39.513517 sequence 0x50 checksum 0x9dcB
May 31 16:47:46.563781 sequence 0x45 checksum 0x7e1a
May 31 16:47:46.563970 sequence 0x46 checksum 0x226d
May 31 16:47:46.564104 sequence 0x52 checksum 0x99cd
May 31 16:47:46.581087 sequence 0x49 checksum 0x8e64
May 31 16:47:46.581222 sequence 0x4e checksum 0x34b3
May 31 16:47:46.581353 sequence 0x50 checksum 0x9dcB
May 31 16:47:55.799090 sequence 0x45 checksum 0x7e1a
May 31 16:47:55.799223 sequence 0x46 checksum 0x226d
May 31 16:47:55.799347 sequence 0x52 checksum 0x99cd
May 31 16:47:55.818255 sequence 0x49 checksum 0x8e64
May 31 16:47:55.818473 sequence 0x4e checksum 0x34b3
May 31 16:47:55.818606 sequence 0x50 checksum 0x9dcB
May 31 16:48:03.455816 sequence 0x49 checksum 0x8e64
May 31 16:48:03.455973 sequence 0x4e checksum 0x34b3
```

**Meaning** The output shows that checksum information is captured in the IS-IS trace log file.

**Related Documentation**

- [Example: Configuring BFD for IS-IS on page 91](#)



## CHAPTER 5

# IS-IS Routing Policy and Route Redistribution

- [Example: Configuring IS-IS Route Leaking From a Level 2 Area to a Level 1 Area on page 55](#)
- [Example: Redistributing OSPF Routes into IS-IS on page 61](#)
- [Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS on page 72](#)
- [Example: Configuring an IS-IS Default Route Policy on Logical Systems on page 82](#)
- [IS-IS Extensions to Support Route Tagging on page 88](#)

### Example: Configuring IS-IS Route Leaking From a Level 2 Area to a Level 1 Area

This example shows how to leak prefixes in an IS-IS network from a Level 2 area to a Level 1 area.

- [Requirements on page 55](#)
- [Overview on page 55](#)
- [Configuration on page 56](#)
- [Verification on page 60](#)

#### Requirements

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

#### Overview

Every routing protocol passes routing information up or down the routing hierarchy. This bidirectional flow of routing information is known as route leaking.


By default, IS-IS protocol leaks routing information from a Level 1 area to a Level 2 area. However, to leak routing information from a Level 2 area to a Level 1 area, an export policy must be explicitly configured.

## Topology

In this example, Devices R3 and R4 are configured in a Level 2 area. Devices R5, R6, and R7 are configured in a Level 1 area.

Figure 10 on page 56 shows the topology used in this example.

Figure 10: Route Leaking From a Level 2 Area to a Level 1 Area



## Configuration

### Configuring Route Leaking From a Level 2 Area to a Level 1 Area

#### CLI Quick Configuration

To quickly configure route leaking from a Level 2 area to a Level 1 area, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device R3
set interfaces fe-1/2/0 unit 0 description to-R4
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.21/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set interfaces lo0 unit 0 family iso address 49.0001.0192.0168.0003.00
set policy-options policy-statement leak-L2-to-L1 from route-filter 192.168.0.0/24 orlonger
set policy-options policy-statement leak-L2-to-L1 from protocol isis
set policy-options policy-statement leak-L2-to-L1 from level 2
set policy-options policy-statement leak-L2-to-L1 to protocol isis
set policy-options policy-statement leak-L2-to-L1 to level 1
set policy-options policy-statement leak-L2-to-L1 then accept
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0 level 1 disable
set protocols isis export leak-L2-to-L1
```

**Device R4**

```

set interfaces fe-1/2/0 unit 0 description to-R3
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.18/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.25/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.0001.0192.0168.0004.00
set policy-options policy-statement leak-L2-to-L1 from route-filter 192.168.0.0/24 orlonger
set policy-options policy-statement leak-L2-to-L1 from protocol isis
set policy-options policy-statement leak-L2-to-L1 from level 2
set policy-options policy-statement leak-L2-to-L1 to protocol isis
set policy-options policy-statement leak-L2-to-L1 to level 1
set policy-options policy-statement leak-L2-to-L1 then accept
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0 level 1 disable
set protocols isis export leak-L2-to-L1

```

**Device R5**

```

set interfaces fe-1/2/0 unit 0 description to-R3
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.22/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R4
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.26/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 description to-R6
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.29/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/3 unit 0 description to-R7
set interfaces fe-1/2/3 unit 0 family inet address 10.0.0.38/30
set interfaces fe-1/2/3 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.5/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0005.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface fe-1/2/2.0 level 2 disable
set protocols isis interface fe-1/2/3.0 level 2 disable
set protocols isis interface lo0.0 level 1 disable

```

**Device R6**

```

set interfaces fe-1/2/0 unit 0 description to-R5
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.30/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R7
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.33/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.6/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0006.00
set protocols isis interface fe-1/2/0.0 level 2 disable
set protocols isis interface fe-1/2/1.0 level 2 disable
set protocols isis interface lo0.0 level 2 disable

```

**Device R7**

```

set interfaces fe-1/2/0 unit 0 description to-R6
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.34/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R5

```

```

set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.37/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.7/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0007.00
set protocols isis interface fe-1/2/0.0 level 2 disable
set protocols isis interface fe-1/2/1.0 level 2 disable
set protocols isis interface lo0.0 level 2 disable

```

### Configuring Route Leaking From a Level 2 Area to a Level 1 Area

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

To configure route leaking from a Level 2 area to a Level 1 area:

1. Configure the network interfaces.

Enable IS-IS on the interfaces by including the ISO address family on each interface.

```

[edit interfaces]
user@R3# set fe-1/2/0 unit 0 description to-R4
user@R3# set fe-1/2/0 unit 0 family inet address 10.0.0.17/30
user@R3# set fe-1/2/0 unit 0 family iso
user@R3# set fe-1/2/1 unit 0 description to-R5
user@R3# set fe-1/2/1 unit 0 family inet address 10.0.0.21/30
user@R3# set fe-1/2/1 unit 0 family iso

```

Similarly, configure other routers.

2. Configure two loopback interface addresses.

One address is for IPv4, and the other address is to enable the router to form adjacencies with other routers in the area.

```

[edit interfaces lo0 unit 0]
user@R3# set family inet address 192.168.0.3/32
user@R3# set family iso address 49.0001.0192.0168.0003.00

```

3. Specify the IS-IS level on a per-interface basis.

```

[edit protocols isis interface]
user@R3# set fe-1/2/0.0 level 1 disable
user@R3# set fe-1/2/1.0 level 1 disable
user@R3# set lo0.0 level 1 disable

```

4. Configure a route leaking policy on the routers configured in the Level 2 area to leak routes into the Level 1 area.

```

[edit policy-options policy-statement leak-L2-to-L1]
user@R3# set from route-filter 192.168.0.0/24 orlonger
user@R3# set from protocol isis
user@R3# set from level 2
user@R3# set to protocol isis
user@R3# set to level 1
user@R3# set then accept

[edit protocols isis]
user@R3# set export leak-L2-to-L1

```



Similarly, configure Router R4.

## Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols isis**, and **show policy-options** commands.

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/2/0 {
 unit 0 {
 description to-R4;
 family inet {
 address 10.0.0.17/30;
 }
 family iso;
 }
}
fe-1/2/1 {
 unit 0 {
 description to-R5;
 family inet {
 address 10.0.0.21/30;
 }
 family iso;
 }
}
lo0 {
 unit 0 {
 family inet {
 address 192.168.0.3/32;
 }
 family iso {
 address 49.0001.0192.0168.0003.00;
 }
 }
}

user@R3# show protocols isis
export leak-L2-to-L1;
interface fe-1/2/0.0 {
 level 1 disable;
}
interface fe-1/2/1.0 {
 level 1 disable;
}
interface lo0.0 {
 level 1 disable;
}

user@R3# show policy-options
policy-statement leak-L2-to-L1 {
 from {
 protocol isis;
 level 2;
 route-filter 192.168.0.0/24 orlonger;
 }
}
```

```

 }
 to {
 protocol isis;
 level 1;
 }
 then accept;
}

```

Similarly, confirm the configuration on all other routers. If you are done configuring the routers, enter **commit** from configuration mode.

## Verification

### Verifying Route Leaking From a Level 2 Area to a Level 1 Area

**Purpose** Verify that IS-IS leaks routes from a Level 2 area to a Level 1 area.

**Action** To verify that route leaking is taking place, use the following commands:

- **show isis adjacency** (to verify that the IS-IS network is up and adjacencies have been established)
- **show isis database detail** (to verify the presence of leaked routes)

1. From operational mode on Router R3, run the **show isis adjacency** command.

```

user@R3> show isis adjacency
Interface System L State Hold (secs) SNPA
fe-1/2/0.0 R4 2 Up 7 0:5:85:8f:94:bd
fe-1/2/1.0 R5 2 Up 7 0:5:85:8f:94:bd

```

The output verifies that the interfaces on Router R3 are up and have established adjacencies with the connecting interfaces on Routers R4 and R5. If you don't see the interfaces being functional, see the [“Results” on page 59](#) section for troubleshooting your configuration.

2. From operational mode on Router R3, run the **show isis database detail** command.

```

user@R3> show isis database detail
IS-IS level 1 link-state database:

R3.00-00 Sequence: 0x19, Checksum: 0x3453, Lifetime: 1078 secs
 IP prefix: 192.168.0.4/32 Metric: 10 Internal Down
 IP prefix: 192.168.0.5/32 Metric: 10 Internal Down
 IP prefix: 192.168.0.6/32 Metric: 20 Internal Down
 IP prefix: 192.168.0.7/32 Metric: 20 Internal Down

IS-IS level 2 link-state database:

R3.00-00 Sequence: 0x1c, Checksum: 0xc657, Lifetime: 1078 secs
 IS neighbor: R4.02 Metric: 10
 IS neighbor: R5.02 Metric: 10
 IP prefix: 10.0.0.16/30 Metric: 10 Internal Up
 IP prefix: 10.0.0.20/30 Metric: 10 Internal Up
 IP prefix: 192.168.0.3/32 Metric: 0 Internal Up

R4.00-00 Sequence: 0x19, Checksum: 0xea13, Lifetime: 1076 secs
 IS neighbor: R4.02 Metric: 10
 IS neighbor: R5.03 Metric: 10

```

```

IP prefix: 10.0.0.16/30 Metric: 10 Internal Up
IP prefix: 10.0.0.24/30 Metric: 10 Internal Up
IP prefix: 192.168.0.4/32 Metric: 0 Internal Up

R4.02-00 Sequence: 0x17, Checksum: 0xecab, Lifetime: 1076 secs
IS neighbor: R3.00 Metric: 0
IS neighbor: R4.00 Metric: 0

R5.00-00 Sequence: 0x12, Checksum: 0xf4e5, Lifetime: 1076 secs
IS neighbor: R5.02 Metric: 10
IS neighbor: R5.03 Metric: 10
IP prefix: 10.0.0.20/30 Metric: 10 Internal Up
IP prefix: 10.0.0.24/30 Metric: 10 Internal Up
IP prefix: 10.0.0.28/30 Metric: 10 Internal Up
IP prefix: 10.0.0.32/30 Metric: 20 Internal Up
IP prefix: 10.0.0.36/30 Metric: 10 Internal Up
IP prefix: 192.168.0.5/32 Metric: 0 Internal Up
IP prefix: 192.168.0.6/32 Metric: 10 Internal Up
IP prefix: 192.168.0.7/32 Metric: 10 Internal Up

R5.02-00 Sequence: 0xb, Checksum: 0x2d74, Lifetime: 1076 secs
IS neighbor: R3.00 Metric: 0
IS neighbor: R5.00 Metric: 0

R5.03-00 Sequence: 0xb, Checksum: 0x6c32, Lifetime: 1076 secs
IS neighbor: R4.00 Metric: 0
IS neighbor: R5.00 Metric: 0

```

The **Down** keyword identifies the routes that have successfully leaked from the Level 2 area to the Level 1 area.

**Meaning** Route leaking from a Level 2 to a Level 1 area is functioning as expected.

## Example: Redistributing OSPF Routes into IS-IS

- [Understanding Routing Policies on page 61](#)
- [Example: Redistributing OSPF Routes into IS-IS on page 64](#)

### Understanding Routing Policies

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

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

- [Importing and Exporting Routes on page 62](#)
- [Active and Inactive Routes on page 63](#)

- [Explicitly Configured Routes on page 63](#)
- [Dynamic Database on page 64](#)

### Importing and Exporting Routes

Two terms—*import* and *export*—explain how routes move between the routing protocols and the routing table.


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



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

As shown in [Figure 11 on page 62](#), you use import routing policies to control which routes are placed in the routing table, and export routing policies to control which routes are advertised from the routing table to neighbors.

**Figure 11: Importing and Exporting Routes**



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

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

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

- Control which routes a routing protocol places in the routing table.
- Control which active routes a routing protocol advertises from the routing table. An *active route* is a route that is chosen from all routes in the routing table to reach a destination.
- Manipulate the route characteristics as a routing protocol places the route in the routing table or advertises the route from the routing table.

You can manipulate the route characteristics to control which route is selected as the active route to reach a destination. The active route is placed in the forwarding table and is used to forward traffic toward the route's destination. In general, the active route is also advertised to a router's neighbors.

### Active and Inactive Routes

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

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

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



**NOTE:** By default, BGP advertises active routes. However, you can configure BGP to advertise *inactive routes*, which go to the same destination as other routes but have less preferable metrics.

### Explicitly Configured Routes

An *explicitly configured route* is a route that you have configured. *Direct routes* are not explicitly configured. They are created as a result of IP addresses being configured on an interface. Explicitly configured routes include aggregate, generated, local, and static routes. (An *aggregate route* is a route that distills groups of routes with common addresses into one route. A *generated route* is a route used when the routing table has no information about how to reach a particular destination. A *local route* is an IP address assigned to a router interface. A *static route* is an unchanging route to a destination.)

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

static routes can be exported from the routing table to routing protocols, whereas local routes cannot.

### Dynamic Database

---

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

### Example: Redistributing OSPF Routes into IS-IS

This example shows how to redistribute OSPF routes into an IS-IS network.

- [Requirements on page 64](#)
- [Overview on page 64](#)
- [Configuration on page 65](#)
- [Verification on page 70](#)

#### Requirements

---

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

#### Overview

---

Export policy can be applied to IS-IS to facilitate route redistribution.


Junos OS does not support the application of import policy for link-state routing protocols like IS-IS because such policies can lead to inconsistent link-state database (LSDB) entries, which in turn can result in routing inconsistencies.

In this example, OSPF routes 192.168.0/24 through 192.168.3/24 are redistributed into IS-IS area 49.0002 from Device R2.

In addition, policies are configured to ensure that Device R1 can reach destinations on the 10.0.0.44/30 network, and that Device R3 can reach destinations on the 10.0.0.36/30 network. This enables end-to-end reachability.

[Figure 12 on page 65](#) shows the topology used in this example.

Figure 12: IS-IS Route Redistribution Topology



“CLI Quick Configuration” on page 65 shows the configuration for all of the devices in Figure 12 on page 65. The section “Step-by-Step Procedure” on page 66 describes the steps on Device R2. “Step-by-Step Procedure” on page 67 describes the steps on Device R3.

### Configuration

#### CLI Quick Configuration

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

**Device R1**

```
set interfaces fe-1/2/0 unit 0 description to-R7
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.38/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 172.16.3.5/32
set interfaces lo0 unit 0 family iso address 49.0002.0172.0016.0305.00
set protocols isis interface fe-1/2/0.38
set protocols isis interface lo0.0
```

**Device R2**

```
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.37/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/0 unit 0 description to-OSPF-network
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.45/30
set interfaces lo0 unit 0 family inet address 172.16.9.7/32
set interfaces lo0 unit 0 family iso address 49.0002.0172.0016.0907.00
set protocols isis export ospf-isis
set protocols isis export send-direct-to-isis-neighbors
set protocols isis interface fe-1/2/1.0
set protocols isis interface lo0.0
```

```

set protocols ospf export send-direct-to-ospf-neighbors
set protocols ospf area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf area 0.0.0.1 interface lo0.0 passive
set policy-options policy-statement ospf-isis term 1 from protocol ospf
set policy-options policy-statement ospf-isis term 1 from route-filter 192.168.0.0/22
 longer
set policy-options policy-statement ospf-isis term 1 then accept
set policy-options policy-statement send-direct-to-isis-neighbors from protocol direct
set policy-options policy-statement send-direct-to-isis-neighbors from route-filter
 10.0.0.44/30 exact
set policy-options policy-statement send-direct-to-isis-neighbors then accept
set policy-options policy-statement send-direct-to-ospf-neighbors from protocol direct
set policy-options policy-statement send-direct-to-ospf-neighbors from route-filter
 10.0.0.36/30 exact
set policy-options policy-statement send-direct-to-ospf-neighbors then accept

```

**Device R3**

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.46/30
set interfaces lo0 unit 0 family inet address 192.168.1.1/32
set interfaces lo0 unit 0 family inet address 192.168.2.1/32
set interfaces lo0 unit 0 family inet address 192.168.3.1/32
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set protocols ospf export ospf
set protocols ospf area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf area 0.0.0.1 interface lo0.0 passive
set policy-options policy-statement ospf term 1 from protocol static
set policy-options policy-statement ospf term 1 then accept
set routing-options static route 192.168.0.0/24 discard
set routing-options static route 192.168.1.0/24 discard
set routing-options static route 192.168.3.0/24 discard
set routing-options static route 192.168.2.0/24 discard

```

#### Step-by-Step Procedure

To configure Device R2:

1. Configure the network interfaces.
 

```

[edit interfaces]
user@R2# set fe-1/2/1 unit 0 description to-R5
user@R2# set fe-1/2/1 unit 0 family inet address 10.0.0.37/30
user@R2# set fe-1/2/1 unit 0 family iso
user@R2# set fe-1/2/0 unit 0 description to-OSPF-network
user@R2# set fe-1/2/0 unit 0 family inet address 10.0.0.45/30
user@R2# set lo0 unit 0 family inet address 172.16.9.7/32
user@R2# set lo0 unit 0 family iso address 49.0002.0172.0016.0907.00

```
2. Configure IS-IS on the interface facing Device R1 and the loopback interface.
 

```

[edit protocols isis]
user@R2# set interface fe-1/2/1.0
user@R2# set interface lo0.0

```
3. Configure the policy that enables Device R1 to reach the 10.0.0.44/30 network.
 

```

[edit policy-options policy-statement send-direct-to-isis-neighbors]
user@R2# set from protocol direct
user@R2# set from route-filter 10.0.0.44/30 exact
user@R2# set then accept

```
4. Apply the policy that enables Device R1 to reach the 10.0.0.44/30 network.



- ```
[edit protocols isis]
user@R2# set export send-direct-to-isis-neighbors
```
5. Configure OSPF on the interfaces.


```
[edit protocols ospf]
user@R2# set area 0.0.0.1 interface fe-1/2/0.0
user@R2# set area 0.0.0.1 interface lo0.0 passive
```
 6. Configure the OSPF route redistribution policy.


```
[edit policy-options policy-statement ospf-isis term 1]
user@R2# set from protocol ospf
user@R2# set from route-filter 192.168.0.0/22 longer
user@R2# set then accept
```
 7. Apply the OSPF route redistribution policy to the IS-IS instance.


```
[edit protocols isis]
user@R2# set export ospf-isis
```
 8. Configure the policy that enables Device R3 to reach the 10.0.0.36/30 network.


```
[edit policy-options policy-statement send-direct-to-ospf-neighbors]
user@R2# set from protocol direct
user@R2# set from route-filter 10.0.0.36/30 exact
user@R2# set then accept
```
 9. Apply the policy that enables Device R3 to reach the 10.0.0.36/30 network.


```
[edit protocols ospf]
user@R2# set export send-direct-to-ospf-neighbors
```

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

To configure multi-level IS-IS:

1. Configure the network interfaces.

Multiple addresses are configured on the loopback interface to simulate multiple route destinations.

```
[edit interfaces]
user@R3# set fe-1/2/0 unit 0 family inet address 10.0.0.46/30
user@R3# set lo0 unit 0 family inet address 192.168.1.1/32
user@R3# set lo0 unit 0 family inet address 192.168.2.1/32
user@R3# set lo0 unit 0 family inet address 192.168.3.1/32
user@R3# set lo0 unit 0 family inet address 192.168.0.1/32
```

2. Configure static routes to the loopback interface addresses.

These are the routes that are redistributed into IS-IS.

```
[edit routing-options static]
user@R3# set route 192.168.0.0/24 discard
user@R3# set route 192.168.1.0/24 discard
user@R3# set route 192.168.3.0/24 discard
user@R3# set route 192.168.2.0/24 discard
```

3. Configure OSPF on the interfaces.

```
[edit protocols ospf area 0.0.0.1]
user@R3# set interface fe-1/2/0.0
user@R3# set interface lo0.0 passive
```

4. Configure the OSPF policy to export the static routes.

```
[edit policy-options policy-statement ospf term 1]
user@R3# set from protocol static
user@R3# set then accept
```

5. Apply the OSPF export policy.

```
[edit protocols ospf]
user@R3# set export ospf
```

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

```
Device R2 user@R2# show interfaces
fe-1/2/1 {
  unit 0 {
    description to-R5;
    family inet {
      address 10.0.0.37/30;
    }
    family iso;
  }
}
fe-1/2/0 {
  unit 0 {
    description to-OSPF-network;
    family inet {
      address 10.0.0.45/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 172.16.9.7/32;
    }
    family iso {
      address 49.0002.0172.0016.0907.00;
    }
  }
}

user@R2# show protocols
isis {
  export [ ospf-isis send-direct-to-isis-neighbors ];
  interface fe-1/2/1.0;
  interface lo0.0;
}
```

```

ospf {
  export send-direct-to-ospf-neighbors;
  area 0.0.0.1 {
    interface fe-1/2/0.0;
    interface lo0.0 {
      passive;
    }
  }
}

user@R2# show policy-options
policy-statement ospf-isis {
  term 1 {
    from {
      protocol ospf;
      route-filter 192.168.0.0/22 longer;
    }
    then accept;
  }
}

policy-statement send-direct-to-isis-neighbors {
  from {
    protocol direct;
    route-filter 10.0.0.44/30 exact;
  }
  then accept;
}

policy-statement send-direct-to-ospf-neighbors {
  from {
    protocol direct;
    route-filter 10.0.0.36/30 exact;
  }
  then accept;
}

```

Device R3

```

user@R3# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.46/30;
    }
  }
}

lo0 {
  unit 0 {
    family inet {
      address 192.168.1.1/32;
      address 192.168.2.1/32;
      address 192.168.3.1/32;
      address 192.168.0.1/32;
    }
  }
}

user@R3# show protocols
ospf {

```

```

export ospf;
area 0.0.0.1 {
    interface fe-1/2/0.0;
    interface lo0.0 {
        passive;
    }
}
}

user@R3# show policy-options
policy-statement ospf {
    term 1 {
        from protocol static;
        then accept;
    }
}

user@R3# show routing-options
static {
    route 192.168.0.0/24 discard;
    route 192.168.1.0/24 discard;
    route 192.168.3.0/24 discard;
    route 192.168.2.0/24 discard;
}

```

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

Verification

Confirm that the configuration is working properly.

- [Verifying OSPF Route Advertisement on page 70](#)
- [Verifying Route Redistribution on page 71](#)
- [Verifying Connectivity on page 71](#)

Verifying OSPF Route Advertisement

Purpose Make sure that the expected routes are advertised by OSPF.

Action From operational mode on Device R2, enter the **show route protocol ospf** command.

```
user@R2> show route protocol ospf
```

```
inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

192.168.0.0/24      *[OSPF/150] 03:54:21, metric 0, tag 0
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.0.1/32     *[OSPF/10] 03:54:21, metric 1
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.1.0/24     *[OSPF/150] 03:54:21, metric 0, tag 0
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.1.1/32     *[OSPF/10] 03:54:21, metric 1
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.2.0/24     *[OSPF/150] 03:54:21, metric 0, tag 0
                   > to 10.0.0.46 via fe-1/2/0.0

```

```

192.168.2.1/32      *[OSPF/10] 03:54:21, metric 1
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.3.0/24     *[OSPF/150] 03:54:21, metric 0, tag 0
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.3.1/32     *[OSPF/10] 03:54:21, metric 1
                   > to 10.0.0.46 via fe-1/2/0.0
224.0.0.5/32       *[OSPF/10] 03:56:03, metric 1
                   MultiRecv

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

Meaning The 192.168/16 routes are advertised by OSPF.

Verifying Route Redistribution

Purpose Make sure that the expected routes are redistributed from OSPF into IS-IS.

Action From operational mode on Device R1, enter the **show route protocol isis** command.

```
user@R1> show route protocol isis
```

```

inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.44/30        *[IS-IS/160] 03:45:24, metric 20
                   > to 10.0.0.37 via fe-1/2/0.0
172.16.9.7/32       *[IS-IS/15] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.0.0/24      *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.0.1/32      *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.1.0/24      *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.1.1/32      *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.2.0/24      *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.2.1/32      *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.3.0/24      *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.3.1/32      *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

Meaning The 192.168/16 routes are redistributed into IS-IS.

Verifying Connectivity

Purpose Check that Device R1 can reach the destinations on Device R3.

Action From operational mode, enter the **ping** command.

```
user@R1> ping 192.168.1.1
```

```
PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=63 time=2.089 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=63 time=1.270 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=63 time=2.135 ms
```

Meaning These results confirm that Device R1 can reach the destinations in the OSPF network.

Related Documentation

- [Example: Configuring IS-IS on page 13](#)

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

- [Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions on page 72](#)
- [Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS on page 73](#)

Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions

A *BGP community* is a group of destinations that share a common property. Community information is included as a path attribute in BGP update messages. This information identifies community members and enables you to perform actions on a group without having to elaborate upon each member. You can use community and extended communities attributes to trigger routing decisions, such as acceptance, rejection, preference, or redistribution.

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

A community value is a 32-bit field that is divided into two main sections. The first 16 bits of the value encode the AS number of the network that originated the community, while the last 16 bits carry a unique number assigned by the AS. This system attempts to guarantee a globally unique set of community values for each AS in the Internet. Junos OS uses a notation of **as-number:community-value**, where each value is a decimal number. The AS values of 0 and 65,535 are reserved, as are all of the community values within those AS numbers. Each community, or set of communities, is given a name within the **[edit policy-options]** configuration hierarchy. The name of the community uniquely identifies it to the routing device and serves as the method by which routes are categorized. For example, a route with a community value of 64510:1111 might belong to the community named **AS64510-routes**. The community name is also used within a routing policy as a match criterion or as an action. The command syntax for creating a community is: **policy-options community name members [community-ids]**. The **community-ids** are either a single community value or multiple community values. When more than one value is assigned to a community name, the routing device interprets this as a logical AND of the community values. In other words, a route must have all of the configured values before being assigned the community name.

The regular community attribute is four octets. Networking enhancements, such as VPNs, have functionality requirements that can be satisfied by an attribute such as a community. However, the 4-octet community value does not provide enough expansion and flexibility to accommodate VPN requirements. This leads to the creation of extended communities. An extended community is an 8-octet value that is also divided into two main sections. The first 2 octets of the community encode a type field while the last 6 octets carry a unique set of data in a format defined by the type field. Extended communities provide a larger range for grouping or categorizing communities.

The BGP extended communities attribute format has three fields:

type:administrator:assigned-number. The routing device expects you to use the words **target** or **origin** to represent the type field. The administrator field uses a decimal number for the AS or an IPv4 address, while the assigned number field expects a decimal number no larger than the size of the field (65,535 for 2 octets or 4,294,967,295 for 4 octets).

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

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

This example defines a policy that takes BGP routes from the **Edu** community and places them into IS-IS with a metric of 63.

- [Requirements on page 73](#)
- [Overview on page 73](#)
- [Configuration on page 74](#)
- [Verification on page 81](#)


Requirements

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

Overview

[Figure 13 on page 74](#) shows the topology used in this example.

Figure 13: Redistributing BGP Routes with a Specific Community Tag into IS-IS



In this example, Device A, Device B, Device C, and Device D are in autonomous system (AS) 1 and are running IS-IS. All of the AS 1 devices, except Device D, are running internal BGP (IBGP).

Device E is in AS 2 and has an external BGP (EBGP) peering session with Device C. Device E has two static routes, 10.2.0.0/16 and 10.3.0.0/16. These routes are tagged with the Edu 2:5 community attribute and are advertised by way of EBGP to Device C.

Device C accepts the BGP routes that are tagged with the Edu 2:5 community attribute, redistributes the routes into IS-IS, and applies an IS-IS metric of 63 to these routes.

“CLI Quick Configuration” on page 74 shows the configuration for all of the devices in Figure 13 on page 74. The section “Step-by-Step Procedure” on page 76 describes the steps on Device C and Device E.

Configuration

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

Device A

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.5/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set protocols bgp group int type internal
set protocols bgp group int local-address 192.168.0.1
set protocols bgp group int neighbor 192.168.0.2
set protocols bgp group int neighbor 192.168.0.3
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface lo0.0
set routing-options router-id 192.168.0.1
set routing-options autonomous-system 1

```

Device B

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.6/30
set interfaces fe-1/2/0 unit 0 family iso

```



```

set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.9/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
set protocols bgp group int type internal
set protocols bgp group int local-address 192.168.0.2
set protocols bgp group int neighbor 192.168.0.1
set protocols bgp group int neighbor 192.168.0.3
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0
set routing-options router-id 192.168.0.2
set routing-options autonomous-system 1

```

```

Device C
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.10/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.13/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.25/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0003.00
set protocols bgp group int type internal
set protocols bgp group int local-address 192.168.0.3
set protocols bgp group int neighbor 192.168.0.1
set protocols bgp group int neighbor 192.168.0.2
set protocols bgp group external-peers type external
set protocols bgp group external-peers export send-isis-and-direct
set protocols bgp group external-peers peer-as 2
set protocols bgp group external-peers neighbor 10.0.0.26
set protocols isis export Edu-to-isis
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface fe-1/2/2.0 level 1 disable
set protocols isis interface fe-1/2/2.0 level 2 passive
set protocols isis interface lo0.0
set policy-options policy-statement Edu-to-isis term 1 from protocol bgp
set policy-options policy-statement Edu-to-isis term 1 from community Edu
set policy-options policy-statement Edu-to-isis term 1 then metric 63
set policy-options policy-statement Edu-to-isis term 1 then accept
set policy-options policy-statement send-isis-and-direct term 1 from protocol isis
set policy-options policy-statement send-isis-and-direct term 1 from protocol direct
set policy-options policy-statement send-isis-and-direct term 1 from route-filter
  10.0.0.0/16 orlonger
set policy-options policy-statement send-isis-and-direct term 1 from route-filter
  192.168.0.0/16 orlonger
set policy-options policy-statement send-isis-and-direct term 1 then accept
set policy-options community Edu members 2:5
set routing-options router-id 192.168.0.3
set routing-options autonomous-system 1

```

```

Device D
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.14/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0004.00

```

```
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface lo0.0
set routing-options router-id 192.168.0.4
set routing-options autonomous-system 1
```

Device E

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.26/30
set interfaces lo0 unit 7 family inet address 192.168.0.5/32 primary
set interfaces lo0 unit 7 family inet address 10.2.0.1/32
set interfaces lo0 unit 7 family inet address 10.3.0.1/32
set protocols bgp group external-peers type external
set protocols bgp group external-peers export statics
set protocols bgp group external-peers peer-as 1
set protocols bgp group external-peers neighbor 10.0.0.25
set policy-options policy-statement statics from protocol static
set policy-options policy-statement statics then community add Edu
set policy-options policy-statement statics then accept
set policy-options community Edu members 2:5
set routing-options static route 10.2.0.0/16 reject
set routing-options static route 10.2.0.0/16 install
set routing-options static route 10.3.0.0/16 reject
set routing-options static route 10.3.0.0/16 install
set routing-options router-id 192.168.0.5
set routing-options autonomous-system 2
```

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

To configure Device E:

1. Configure the interfaces.

```
[edit interfaces]
user@E# set fe-1/2/0 unit 0 family inet address 10.0.0.26/30
user@E# set lo0 unit 7 family inet address 192.168.0.5/32 primary
user@E# set lo0 unit 7 family inet address 10.2.0.1/32
user@E# set lo0 unit 7 family inet address 10.3.0.1/32
```

2. Configure the **statics** policy, which adds the **Edu** community attribute to the static routes.

```
[edit policy-options]
user@E# set policy-statement statics from protocol static
user@E# set policy-statement statics then community add Edu
user@E# set policy-statement statics then accept
user@E# set community Edu members 2:5
```

3. Configure EBGp and apply the **statics** policy.

```
[edit protocols bgp group external-peers]
user@E# set type external
user@E# set export statics
user@E# set peer-as 1
user@E# set protocols bgp group external-peers neighbor 10.0.0.25
```

4. Configure the static routes.

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

```

user@E# set route 10.2.0.0/16 reject
user@E# set route 10.2.0.0/16 install
user@E# set route 10.3.0.0/16 reject
user@E# set route 10.3.0.0/16 install

```

5. Configure the router ID and the AS number.

```

[edit routing-options]
user@E# set router-id 192.168.0.5
user@E# set autonomous-system 2

```

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

To configure Device C:

1. Configure the interfaces.

```

[edit interfaces]
user@C# set fe-1/2/0 unit 0 family inet address 10.0.0.10/30
user@C# set fe-1/2/0 unit 0 family iso
user@C# set fe-1/2/1 unit 0 family inet address 10.0.0.13/30
user@C# set fe-1/2/1 unit 0 family iso
user@C# set fe-1/2/2 unit 0 family inet address 10.0.0.25/30
user@C# set fe-1/2/2 unit 0 family iso
user@C# set lo0 unit 0 family inet address 192.168.0.3/32
user@C# set lo0 unit 0 family iso address 49.0002.0192.0168.0003.00

```

2. Configure IBGP.

```

[edit protocols bgp group int]
user@C# set type internal
user@C# set local-address 192.168.0.3
user@C# set neighbor 192.168.0.1
user@C# set neighbor 192.168.0.2

```

3. Configure the Edu-to-isis policy, which redistributes the Edu-tagged BGP routes learned from Device E and applies a metric of 63.

```

[edit policy-options]
user@C# set policy-statement Edu-to-isis term 1 from protocol bgp
user@C# set policy-statement Edu-to-isis term 1 from community Edu
user@C# set policy-statement Edu-to-isis term 1 then metric 63
user@C# set policy-statement Edu-to-isis term 1 then accept
user@C# set community Edu members 2:5

```

4. Enable IS-IS on the interfaces, and apply the Edu-to-isis policy.

```

[edit protocols isis]
user@C# set export Edu-to-isis
user@C# set interface fe-1/2/0.0 level 1 disable
user@C# set interface fe-1/2/1.0 level 1 disable
user@C# set interface fe-1/2/2.0 level 1 disable
user@C# set interface fe-1/2/2.0 level 2 passive
user@C# set interface lo0.0

```

5. Configure the send-isis-and-direct policy, which redistributes routes to Device E, through EBGp.

Without this policy, Device E would not have connectivity to the networks in AS 1.

```
[edit policy-options policy-statement send-isis-and-direct term 1]
user@C# set from protocol isis
user@C# set from protocol direct
user@C# set from route-filter 10.0.0.0/16 orlonger
user@C# set from route-filter 192.168.0.0/16 orlonger
user@C# set then accept
```

6. Configure EBGp and apply the send-isis-and-direct policy.

```
[edit protocols bgp group external-peers]
user@C# set type external
user@C# set export send-isis-and-direct
user@C# set peer-as 2
user@C# set neighbor 10.0.0.26
```

7. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@C# set router-id 192.168.0.3
user@C# set autonomous-system 1
```

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

```
Device E user@E# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.26/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.5/32 {
        primary;
      }
      address 10.2.0.1/32;
      address 10.3.0.1/32;
    }
  }
}

user@E# show protocols
bgp {
  group external-peers {
    type external;
    export statics;
    peer-as 1;
```

```

        neighbor 10.0.0.25;
    }
}

user@E# show policy-options
policy-statement statics {
    from protocol static;
    then {
        community add Edu;
        accept;
    }
}
community Edu members 2:5;

user@E# show routing-options
static {
    route 10.2.0.0/16 {
        reject;
        install;
    }
    route 10.3.0.0/16 {
        reject;
        install;
    }
}
router-id 192.168.0.5;
autonomous-system 2;

Device C user@C# show interfaces
fe-1/2/0 {
    unit 0 {
        family inet {
            address 10.0.0.10/30;
        }
        family iso;
    }
}
fe-1/2/1 {
    unit 0 {
        family inet {
            address 10.0.0.13/30;
        }
        family iso;
    }
}
fe-1/2/2 {
    unit 0 {
        family inet {
            address 10.0.0.25/30;
        }
        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.3/32;
        }
    }
}

```

```
    }  
    family iso {  
        address 49.0002.0192.0168.0003.00;  
    }  
}  
}
```

user@C# show protocols

```
bgp {  
    group int {  
        type internal;  
        local-address 192.168.0.3;  
        neighbor 192.168.0.1;  
        neighbor 192.168.0.2;  
    }  
    group external-peers {  
        type external;  
        export send-isis-and-direct;  
        peer-as 2;  
        neighbor 10.0.0.26;  
    }  
}  
isis {  
    export Edu-to-isis;  
    interface fe-1/2/0.0 {  
        level 1 disable;  
    }  
    interface fe-1/2/1.0 {  
        level 1 disable;  
    }  
    interface fe-1/2/2.0 {  
        level 1 disable;  
        level 2 passive;  
    }  
    interface lo0.0;  
}
```

user@C# show policy-options

```
policy-statement Edu-to-isis {  
    term 1 {  
        from {  
            protocol bgp;  
            community Edu;  
        }  
        then {  
            metric 63;  
            accept;  
        }  
    }  
}  
policy-statement send-isis-and-direct {  
    term 1 {  
        from {  
            protocol [ isis direct ];  
            route-filter 10.0.0.0/16 orlonger;  
            route-filter 192.168.0.0/16 orlonger;
```

```

    }
    then accept;
  }
}
community Edu members 2:5;

user@C# show routing-options
router-id 192.168.0.3;
autonomous-system 1;

```

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

Verification

Confirm that the configuration is working properly.

Verifying the IS-IS Neighbor

Purpose Verify that the BGP routes from Device E are communicated on the IS-IS network in AS 1.

Action From operational mode, enter the **show route protocol isis** command.

```

user@D> show route protocol isis
inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30      *[IS-IS/18] 22:30:53, metric 30
                 > to 10.0.0.13 via fe-1/2/0.0
10.0.0.8/30      *[IS-IS/18] 22:30:53, metric 20
                 > to 10.0.0.13 via fe-1/2/0.0
10.0.0.24/30     *[IS-IS/18] 03:31:21, metric 20
                 > to 10.0.0.13 via fe-1/2/0.0
10.2.0.0/16     *[IS-IS/165] 02:36:31, metric 73
                 > to 10.0.0.13 via fe-1/2/0.0
10.3.0.0/16     *[IS-IS/165] 02:36:31, metric 73
                 > to 10.0.0.13 via fe-1/2/0.0
192.168.0.1/32   *[IS-IS/18] 03:40:28, metric 30
                 > to 10.0.0.13 via fe-1/2/0.0
192.168.0.2/32   *[IS-IS/18] 22:30:53, metric 20
                 > to 10.0.0.13 via fe-1/2/0.0
192.168.0.3/32   *[IS-IS/18] 22:30:53, metric 10
                 > to 10.0.0.13 via fe-1/2/0.0

```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Meaning As expected, the 10.2.0.0/16 and 10.3.0.0/16 routes are in Device D's routing table as IS-IS external routes with a metric of 73. If Device C had not added 63 to the metric, Device D would have a metric of 10 for these routes.

Related Documentation

- [Example: Redistributing OSPF Routes into IS-IS on page 61](#)

Example: Configuring an IS-IS Default Route Policy on Logical Systems

- [Understanding Default Routes on page 82](#)
- [Example: Configuring an IS-IS Default Route Policy on Logical Systems on page 82](#)

Understanding Default Routes

A default route is the route that takes effect when no other route is available for an IP destination address.

If a packet is received on a routing device, the device first checks to see if the IP destination address is on one of the device's local subnets. If the destination address is not local, the device checks its routing table. If the remote destination subnet is not listed in the routing table, the packet is forwarded to the next hop toward the destination using the default route. The default route generally has a next-hop address of another routing device, which performs the same process. The process repeats until a packet is delivered to the destination.

The route evaluation process in each router uses the longest prefix match method to obtain the most specific route. The network with the longest subnet mask that matches the destination IP address is the next-hop network gateway.

The default route in IPv4 is designated as 0.0.0.0/0 or simply 0/0. Similarly, in IPv6, the default route is specified as ::/0. The subnet mask /0 specifies all networks, and is the shortest match possible. A route lookup that does not match any other route uses this route if it is configured and active in the routing table. To be active, the configured next-hop address must be reachable.

Administrators generally point the default route toward the routing device that has a connection to a network service provider. Therefore, packets with destinations outside the organization's local area network, typically destinations on the Internet or a wide area network, are forwarded to the routing device with the connection to that provider. The device to which the default route points is often called the default gateway.

Example: Configuring an IS-IS Default Route Policy on Logical Systems

This example shows logical systems configured on a single physical router and explains how to configure a default route on one logical system.

- [Requirements on page 82](#)
- [Overview on page 83](#)
- [Configuration on page 83](#)
- [Verification on page 86](#)

Requirements

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


Overview

This example shows a logical system redistributing a default route to other logical systems. All logical systems are running IS-IS. A common reason for a default route is to provide a path for sending traffic destined outside the IS-IS domain.

In this example, the default route is not used for forwarding traffic. The **no-install** statement prevents the route from being installed in the forwarding table of Logical System LS3. 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. The **discard** statement silently drops packets without notice.

Figure 14 on page 83 shows the sample network.

Figure 14: IS-IS Logical Systems with a Default Route to an ISP



Configuration

CLI Quick Configuration

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

```
set logical-systems LS3 interfaces lt-1/2/0 unit 3 description LS3->LS2
set logical-systems LS3 interfaces lt-1/2/0 unit 3 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 3 peer-unit 4
set logical-systems LS3 interfaces lt-1/2/0 unit 3 family inet address 10.0.2.1/30
set logical-systems LS3 interfaces lt-1/2/0 unit 3 family iso
set logical-systems LS3 interfaces lt-1/2/0 unit 5 description LS3->LS1
set logical-systems LS3 interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 5 peer-unit 0
set logical-systems LS3 interfaces lt-1/2/0 unit 5 family inet address 10.0.1.1/30
set logical-systems LS3 interfaces lt-1/2/0 unit 5 family iso
```

```

set logical-systems LS3 interfaces lo0 unit 3 family iso address 49.0001.1234.1600.2231.00
set logical-systems LS3 protocols isis export isis-default
set logical-systems LS3 protocols isis interface lt-1/2/0.3
set logical-systems LS3 protocols isis interface lt-1/2/0.5
set logical-systems LS3 protocols isis interface lo0.3 passive
set logical-systems LS3 routing-options static route 0.0.0.0/0 discard
set logical-systems LS3 routing-options static route 0.0.0.0/0 no-install
set logical-systems LS3 policy-options policy-statement isis-default from protocol static
set logical-systems LS3 policy-options policy-statement isis-default from route-filter
  0.0.0.0/0 exact
set logical-systems LS3 policy-options policy-statement isis-default then accept

```

Step-by-Step Procedure

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

To configure an IS-IS default route policy on logical systems:

1. Configure the logical tunnel interfaces.


```

[edit logical-systems LS3 interfaces lt-1/2/0]
user@R1# set unit 3 description LS3->LS2
user@R1# set unit 3 encapsulation ethernet
user@R1# set unit 3 peer-unit 4
user@R1# set unit 3 family inet address 10.0.2.1/30
user@R1# set unit 3 family iso
user@R1# set unit 5 description LS3->LS1
user@R1# set unit 5 encapsulation ethernet
user@R1# set unit 5 peer-unit 0
user@R1# set unit 5 family inet address 10.0.1.1/30
user@R1# set unit 5 family iso
[edit logical-systems LS3 interfaces lo0 unit 3]
user@R1# set family iso address 49.0001.1234.1600.2231.00

```
2. Enable IS-IS on the interfaces.


```

[edit logical-systems LS3 protocols isis]
user@R1# set interface lt-1/2/0.3
user@R1# set interface lt-1/2/0.5
user@R1# set interface lo0.3 passive

```
3. Configure the default route on Logical System LS3.


```

[edit logical-systems LS3 routing-options]
user@R1# set static route 0.0.0.0/0 discard
user@R1# set static route 0.0.0.0/0 no-install

```
4. Configure the default route policy on Logical System LS3.


```

[edit logical-systems LS3 policy-options]
user@R1# set policy-statement isis-default from protocol static
user@R1# set policy-statement isis-default from route-filter 0.0.0.0/0 exact
user@R1# set policy-statement isis-default then accept

```
5. Apply the export policy to IS-IS on Logical System LS3.


```

[edit logical-systems LS3 protocols isis]
user@R1# set export isis-default

```

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

```
[edit]
user@R1# commit
```

Results

From configuration mode, confirm your configuration by issuing the **show logical-systems LS3** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show logical-systems LS3
interfaces {
  lt-1/2/0 {
    unit 3 {
      description LS3->LS2;
      encapsulation ethernet;
      peer-unit 4;
      family inet {
        address 10.0.2.1/30;
      }
      family iso;
    }
    unit 5 {
      description LS3->LS1;
      encapsulation ethernet;
      peer-unit 0;
      family inet {
        address 10.0.1.1/30;
      }
      family iso;
    }
  }
  lo0 {
    unit 3 {
      family iso {
        address 49.0001.1234.1600.2231.00;
      }
    }
  }
}
protocols {
  isis {
    export isis-default;
    interface lt-1/2/0.3;
    interface lt-1/2/0.5;
    interface lo0.3 {
      passive;
    }
  }
}
policy-options {
  policy-statement isis-default {
    from {
      protocol static;
    }
  }
}
```

```
        route-filter 0.0.0.0/0 exact;
      }
      then accept;
    }
  }
  routing-options {
    static {
      route 0.0.0.0/0 {
        discard;
        no-install;
      }
    }
  }
}
```

Verification

Confirm that the configuration is working properly.

Verifying That the Static Route Is Redistributed

Purpose Make sure that the IS-IS policy is working by checking the routing tables.

```

Action user@R1> show route logical-system LS3
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:00:45
                   Discard
10.0.0.0/30        *[IS-IS/15] 1w0d 10:14:14, metric 20
                   to 10.0.2.2 via lt-1/2/0.3
                   > to 10.0.1.2 via lt-1/2/0.5
10.0.1.0/30        *[Direct/0] 1w0d 10:15:18
                   > via lt-1/2/0.5
10.0.1.1/32        *[Local/0] 1w0d 10:15:18
                   Local via lt-1/2/0.5
10.0.2.0/30        *[Direct/0] 1w0d 10:15:18
                   > via lt-1/2/0.3
10.0.2.1/32        *[Local/0] 1w0d 10:15:18
                   Local via lt-1/2/0.3

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1234.1600.2231/72
                   *[Direct/0] 1w0d 10:17:19
                   > via lo0.3

user@R1> show route logical-system LS2
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[IS-IS/160] 00:01:38, metric 10
                   > to 10.0.2.1 via lt-1/2/0.4
10.0.0.0/30        *[Direct/0] 1w0d 10:16:11
                   > via lt-1/2/0.1
10.0.0.2/32        *[Local/0] 1w0d 10:16:11
                   Local via lt-1/2/0.1
10.0.1.0/30        *[IS-IS/15] 1w0d 10:15:07, metric 20
                   > to 10.0.0.1 via lt-1/2/0.1
                   to 10.0.2.1 via lt-1/2/0.4
10.0.2.0/30        *[Direct/0] 1w0d 10:16:11
                   > via lt-1/2/0.4
10.0.2.2/32        *[Local/0] 1w0d 10:16:11
                   Local via lt-1/2/0.4

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.2002/72
                   *[Direct/0] 1w0d 10:18:12
                   > via lo0.2

user@R1> show route logical-system LS1
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[IS-IS/160] 00:02:01, metric 10
                   > to 10.0.1.1 via lt-1/2/0.0
10.0.0.0/30        *[Direct/0] 1w0d 10:16:34
                   > via lt-1/2/0.2
10.0.0.1/32        *[Local/0] 1w0d 10:16:34
                   Local via lt-1/2/0.2

```

```

10.0.1.0/30      *[Direct/0] 1w0d 10:16:34
                  > via lt-1/2/0.0
10.0.1.2/32      *[Local/0] 1w0d 10:16:34
                  Local via lt-1/2/0.0
10.0.2.0/30      *[IS-IS/15] 1w0d 10:15:55, metric 20
                  to 10.0.1.1 via lt-1/2/0.0
                  > to 10.0.0.2 via lt-1/2/0.2

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.1001/72
                  *[Direct/0] 1w0d 10:18:35
                  > via lo0.1

```

Meaning The routing table on Logical System LS3 contains the default 0.0.0.0/0 route from protocol **Static**. The routing tables on Logical System LS1 and Logical System LS2 contain the default 0.0.0.0/0 route from protocol **IS-IS**. If Logical System LS1 and Logical System LS2 receive packets destined for networks not specified in their routing tables, those packets will be sent to Logical System LS3 for further processing. This configuration assumes that Logical System LS3 has a connection to an ISP or another external network.

Related Documentation

- [Example: Configuring IS-IS on page 13](#)
- [Example: Configuring IS-IS on Logical Systems Within the Same Router on page 28](#)

IS-IS Extensions to Support Route Tagging

To control the transmission of routes into IS-IS, or to control transmission of IS-IS routes between different IS-IS levels, you can tag routes with certain attributes. IS-IS routes can carry these attributes, which the routing policies can use to export and import routes between different IS-IS levels. A sub-TLV to the IP prefix TLV is used to carry the tag or attribute on the routes.



NOTE: Route tagging does not work when IS-IS traffic engineering is disabled.

```

protocols {
  isis {
    export tag-lo0;
  }
}
policy-options {
  policy-statement tag-lo0 {
    from {
      interface lo0.0;
    }
    then {
      accept;
      tag 200;
    }
  }
}

```

You can verify that the tag has been correctly applied by using the **show isis database extensive** command. In the command output, look for the **Administrative tag** field.

After verifying that the routes are tagged correctly, you can apply a route leaking policy to match against the presence of administrative tags, rather than specifying a list of route filters.

```
protocols {
  isis {
    export leak-tagged-L2-to-L1;
  }
}
policy-options {
  policy-statement leak-tagged-L2-to-L1 {
    from {
      tag 200;
      protocol isis;
      level 2;
    }
    to {
      protocol isis;
      level 1;
    }
    then accept;
  }
}
```

Related •
Documentation

CHAPTER 6

IS-IS Bidirectional Forwarding Detection

- [Example: Configuring BFD for IS-IS on page 91](#)
- [Example: Configuring BFD Authentication for IS-IS on page 99](#)

Example: Configuring BFD for IS-IS

- [Understanding BFD for IS-IS on page 91](#)
- [Example: Configuring BFD for IS-IS on page 93](#)

Understanding BFD for IS-IS

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than the failure detection mechanisms of IS-IS, providing faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (RX) interval by two if the local BFD instance is the reason for the session flap. The transmission (TX) interval is increased by two if the remote BFD instance is the reason for the session flap.

You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.



NOTE: BFD for IS-IS on an IPv6-only interface is not supported. However, if the interface is dual-stacked (both IPv4 and IPv6 are configured), then you can configure BFD as a client on the IPv4 IS-IS session.

To detect failures in the network, the following set of statements are used in the configuration.

Table 3: Configuring BFD for IS-IS

Statement	Description
bfd-liveness-detection	Enable failure detection.
minimum-interval milliseconds	<p>Specify the minimum transmit and receive intervals for failure detection.</p> <p>This value represents the minimum interval at which the local router transmits hellos packets as well as the minimum interval at which the router expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately.</p> <p>NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.</p> <p>Depending on your network environment, these additional recommendations might apply:</p> <ul style="list-style-type: none"> For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions. For very large-scale network deployments with a large number of BFD sessions, please contact Juniper Networks customer support for more information. For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.
minimum-receive-interval milliseconds	<p>Specify only the minimum receive interval for failure detection.</p> <p>This value represents the minimum interval at which the local router expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number from 1 through 255,000 milliseconds.</p>
multiplier number	<p>Specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down.</p> <p>The default is 3, and you can configure a value from 1 through 225.</p>
no-adaptation	<p>Disable BFD adaptation.</p> <p>In Junos OS Release 9.0 and later, you can specify that the BFD sessions not adapt to changing network conditions.</p> <p>NOTE: We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.</p>
threshold	<ul style="list-style-type: none"> Specify the threshold for the adaptation of the detection time. <p>When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent.</p> Specify the threshold for the transmit interval. <p>NOTE: The threshold value must be greater than the minimum transmit interval multiplied by the multiplier number.</p>

Table 3: Configuring BFD for IS-IS (*continued*)

Statement	Description
transmit-interval minimum-interval	Specify the minimum transmit interval for failure detection. This value represents the minimum interval at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value from 1 through 255,000 milliseconds.
version	Specify the BFD version used for detection. The default is to have the version detected automatically.



NOTE: You can trace BFD operations by including the `traceoptions` statement at the `[edit protocols bfd]` hierarchy level.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

Example: Configuring BFD for IS-IS

This example describes how to configure the Bidirectional Forwarding Detection (BFD) protocol to detect failures in an IS-IS network.

- [Requirements on page 93](#)
- [Overview on page 93](#)
- [Configuration on page 94](#)
- [Verification on page 97](#)

Requirements

Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 14](#) for information about the required IS-IS configuration.

This example uses the following hardware and software components:


- Junos OS Release 7.3 or later
- M Series, MX Series, and T Series routers

Overview

This example shows two routers connected to each other. A loopback interface is configured on each router. IS-IS and BFD protocols are configured on both routers.

[Figure 15 on page 94](#) shows the sample network.

Figure 15: Configuring BFD for IS-IS



Configuration

CLI Quick Configuration

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

Router R1

```
set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 5
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 2
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 3
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic
```

Router R2

```
set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 6
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 3
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 4
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic
```

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



NOTE: To simply configure BFD for IS-IS, only the `minimum-interval` statement is required. The BFD protocol selects default parameters for all the other configuration statements when you use the `bfd-liveness-detection` statement without specifying any parameters.



NOTE: You can change parameters at any time without stopping or restarting the existing session. BFD automatically adjusts to the new parameter value. However, no changes to BFD parameters take place until the values resynchronize with each BFD peer.

To configure BFD for IS-IS on Routers R1 and R2:

1. Enable BFD failure detection for IS-IS.


```
[edit protocols isis]
user@R1# set interface so-0/0/0 bfd-liveness-detection

[edit protocols isis]
user@R2# set interface so-0/0/0 bfd-liveness-detection
```
2. Configure the threshold for the adaptation of the detection time, which must be greater than the multiplier number multiplied by the minimum interval.


```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set detection-time threshold 5

[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set detection-time threshold 6
```
3. Configure the minimum transmit and receive intervals for failure detection.


```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set minimum-interval 2

[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set minimum-interval 3
```
4. Configure only the minimum receive interval for failure detection.


```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set minimum-receive-interval 1

[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set minimum-receive-interval 1
```
5. Disable BFD adaptation.


```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set no-adaptation
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set no-adaptation
```

6. Configure the threshold for the transmit interval, which must be greater than the minimum transmit interval.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval threshold 3
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval threshold 4
```

7. Configure the minimum transmit interval for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval minimum-interval 1
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval minimum-interval 1
```

8. Configure the multiplier number, which is the number of hello packets not received by the neighbor that causes the originating interface to be declared down.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set multiplier 2
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set multiplier 2
```

9. Configure the BFD version used for detection.

The default is to have the version detected automatically.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set version automatic
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set version automatic
```

Results

From configuration mode, confirm your configuration by issuing the **show protocols isis interface** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols isis interface so-0/0/0
```

```
    bfd-liveness-detection {
      version automatic;
      minimum-interval 2;
      minimum-receive-interval 1;
      multiplier 2;
      no-adaptation;
      transmit-interval {
        minimum-interval 1;
        threshold 3;
      }
      detection-time {
        threshold 5;
      }
    }
...
```

```
user@R2# show protocols isis interface so-0/0/0
```

```
    bfd-liveness-detection {
      version automatic;
      minimum-interval 3;
      minimum-receive-interval 1;
      multiplier 2;
      no-adaptation;
      transmit-interval {
        minimum-interval 1;
        threshold 4;
      }
      detection-time {
        threshold 6;
      }
    }
  ...
```

Verification

Confirm that the configuration is working properly.

- [Verifying the Connection Between Routers R1 and R2 on page 97](#)
- [Verifying That IS-IS Is Configured on page 98](#)
- [Verifying That BFD Is configured on page 98](#)

Verifying the Connection Between Routers R1 and R2

Purpose Make sure that Routers R1 and R2 are connected to each other.

Action Ping the other router to check the connectivity between the two routers as per the network topology.

```
user@R1> ping 10.0.0.2
```

```
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=64 time=1.367 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=1.662 ms
64 bytes from 10.0.0.2: icmp_seq=2 ttl=64 time=1.291 ms
^C
--- 10.0.0.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.291/1.440/1.662/0.160 ms
```

```
user@R2> ping 10.0.0.1
```

```
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=64 time=1.287 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=1.310 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1.289 ms
^C
--- 10.0.0.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.287/1.295/1.310/0.010 ms
```

Meaning Routers R1 and R2 are connected to each other.

Verifying That IS-IS Is Configured

Purpose Make sure that the IS-IS instance is running on both routers.

Action Use the **show isis database** statement to check if the IS-IS instance is running on both routers, R1 and R2.

```
user@R1> show isis database
```

```
IS-IS level 1 link-state database:
LSP ID      Sequence Checksum Lifetime Attributes
R1.00-00    0x4a571  0x30c5    1195 L1 L2
R2.00-00    0x4a586  0x4b7e    1195 L1 L2
R2.02-00    0x330ca1 0x3492    1196 L1 L2
  3 LSPs
```

```
IS-IS level 2 link-state database:
LSP ID      Sequence Checksum Lifetime Attributes
R1.00-00    0x4a856  0x5db0    1194 L1 L2
R2.00-00    0x4a89d  0x149b    1194 L1 L2
R2.02-00    0x1fb2ff 0xd302    1194 L1 L2
  3 LSPs
```

```
user@R2> show isis database
```

```
IS-IS level 1 link-state database:
LSP ID      Sequence Checksum Lifetime Attributes
R1.00-00    0x4b707  0xcc80    1195 L1 L2
R2.00-00    0x4b71b  0xeb37    1198 L1 L2
R2.02-00    0x33c2ce 0xb52d    1198 L1 L2
  3 LSPs
```

```
IS-IS level 2 link-state database:
LSP ID      Sequence Checksum Lifetime Attributes
R1.00-00    0x4b9f2  0xee70    1192 L1 L2
R2.00-00    0x4ba41  0x9862    1197 L1 L2
R2.02-00    0x3      0x6242    1198 L1 L2
  3 LSPs
```

Meaning IS-IS is configured on both routers, R1 and R2.

Verifying That BFD Is configured

Purpose Make sure that the BFD instance is running on both routers, R1 and R2.

Action Use the **show bfd session detail** statement to check if BFD instance is running on the routers.

```
user@R1> show bfd session detail
```

```
Address      State      Interface      Detect   Transmit
10.0.0.2     Up         so-0/0/0       Time    Interval  Multiplier
                2.000      1.000          2
Client ISIS R2, TX interval 0.001, RX interval 0.001
Client ISIS R1, TX interval 0.001, RX interval 0.001
Session down time 00:00:00, previous up time 00:00:15
Local diagnostic NbrSignal, remote diagnostic NbrSignal
Remote state AdminDown, version 1
Router 3, routing table index 17
```



```
1 sessions, 2 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

```
user@R2> show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.1	Up	so-0/0/0	2.000	1.000	2

```
Client ISIS R2, TX interval 0.001, RX interval 0.001
Session down time 00:00:00, previous up time 00:00:05
Local diagnostic NbrSignal, remote diagnostic NbrSignal
Remote state AdminDown, version 1
Router 2, routing table index 15
```

```
1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

Meaning BFD is configured on Routers R1 and R2 for detecting failures in the IS-IS network.

Related Documentation

- [Example: Configuring BFD Authentication for IS-IS on page 99](#)

Example: Configuring BFD Authentication for IS-IS

- [Understanding BFD Authentication for IS-IS on page 99](#)
- [Configuring BFD Authentication for IS-IS on page 101](#)
- [Example: Configuring BFD Authentication for IS-IS on page 104](#)

Understanding BFD Authentication for IS-IS

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when running BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over IS-IS. BFD authentication is only supported in the domestic image and is not available in the export image.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 100](#)
- [Security Authentication Keychains on page 100](#)
- [Strict Versus Loose Authentication on page 101](#)

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords might be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm 1 for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm 1. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



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

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session,

and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Configuring BFD Authentication for IS-IS

Beginning with Junos OS Release 9.6, you can configure authentication for BFD sessions running over IS-IS. Routing instances are also supported. Only three steps are needed to configure authentication on a BFD session:

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

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

- [Configuring BFD Authentication Parameters on page 101](#)
- [Viewing Authentication Information for BFD Sessions on page 102](#)

Configuring BFD Authentication Parameters

To configure BFD authentication:

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

[edit]

```
user@host# set protocols isis interface if1-isis bfd-liveness-detection authentication
algorithm keyed-sha-1
```



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

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

This should match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit]
user@host# set protocols isis interface if1-isis bfd-liveness-detection authentication
keychain bfd-isis
```



NOTE: The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 2.
- At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
user@host# set authentication-key-chains key-chain bfd-sr4 key 53 secret
$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit]
user@host> set protocols isis interface if1-isis bfd-liveness-detection authentication
loose-check
```

5. (Optional) View your configuration using the **show bfd session detail** or **show bfd session extensive** command.

6. Repeat these steps to configure the other end of the BFD session.



NOTE: BFD authentication is only supported in the domestic image and is not available in the export image.

Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the if1-isis interface. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-isis**. The authentication keychain is configured with two keys. Key 1 contains the secret data "**\$9\$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm**" and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data "**\$9\$a5jiKW9L.reP38ny.TszF2/9**" and a start time of June 1, 2009, at 3:29:20 PM PST.

```

[edit protocols isis]
interface ifl-isis {
  bfd-liveness-detection {
    authentication {
      algorithm keyed-sha-1;
      key-chain bfd-isis;
    }
  }
}
[edit security]
authentication key-chains {
  key-chain bfd-isis {
    key 1 {
      secret "$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm";
      start-time "2009-6-1.09:46:02 -0700";
    }
    key 2 {
      secret "$9$a5jiKW9l.reP38ny.TszF2/9";
      start-time "2009-6-1.15:29:20 -0700";
    }
  }
}

```

If you commit these updates to your configuration, you see output similar to the following. In the output for the **show bfd sessions detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the **show bfd sessions extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

show bfd sessions detail

```
user@host# show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.9.1.29	Up	ge-4/0/0.0	0.600	0.200	3

Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3, **Authenticate**

Session up time 3d 00:34, previous down time 00:00:01
 Local diagnostic NbrSignal, remote diagnostic AdminDown
 Remote state Up, version 1

1 sessions, 1 clients
 Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

show bfd sessions extensive

```
user@host# show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.9.1.29	Up	ge-4/0/0.0	0.600	0.200	3

Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3, **Authenticate**

keychain bfd-isis, algo keyed-sha-1, mode strict
 Session up time 00:04:42
 Local diagnostic None, remote diagnostic NbrSignal

```

Remote state Up, version 1
Replicated
Min async interval 0.300, min slow interval 1.000
Adaptive async TX interval 0.300, RX interval 0.300
Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3
Remote min TX interval 0.300, min RX interval 0.300, multiplier 3
Local discriminator 2, remote discriminator 2
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-isis, algo keyed-sha-1, mode strict

1 sessions, 1 clients

Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

```

Example: Configuring BFD Authentication for IS-IS

This example shows how to configure BFD authentication for IS-IS.

- [Requirements on page 104](#)
- [Overview on page 104](#)
- [Configuration on page 104](#)
- [Verification on page 106](#)

Requirements


Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 14](#) for information about the required IS-IS configuration.

Overview

In this example, a BFD authentication keychain is configured with meticulous keyed MD5 authentication.

[Figure 16 on page 104](#) shows the topology used in this example.

Figure 16: IS-IS BFD Authentication Topology



[“CLI Quick Configuration” on page 104](#) shows the configuration for both of the devices in [Figure 16 on page 104](#). The section [“Step-by-Step Procedure” on page 105](#) describes the steps on Device R1.

Configuration

CLI Quick Configuration

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

Device R1

```

set security authentication-key-chains key-chain secret123 description for-isis-bfd
set security authentication-key-chains key-chain secret123 key 1 secret "$9$cW-yrv"
set security authentication-key-chains key-chain secret123 key 1 start-time
  "2012-5-31.13:00:00 -0700"
set security authentication-key-chains key-chain secret123 key 2 secret "$9$m5T3"
set security authentication-key-chains key-chain secret123 key 2 start-time
  "2013-5-31.13:00:00 -0700"
set security authentication-key-chains key-chain secret123 key 3 secret "$9$mTQn"
set security authentication-key-chains key-chain secret123 key 3 start-time
  "2014-5-31.13:00:00 -0700"
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection minimum-interval 100
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication key-chain
  secret123
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication algorithm
  meticulous-keyed-md5

```

Device R2

```

set security authentication-key-chains key-chain secret123 description for-isis-bfd
set security authentication-key-chains key-chain secret123 key 1 secret "$9$cW-yrv"
set security authentication-key-chains key-chain secret123 key 1 start-time
  "2012-5-31.13:00:00 -0700"
set security authentication-key-chains key-chain secret123 key 2 secret "$9$m5T3"
set security authentication-key-chains key-chain secret123 key 2 start-time
  "2013-5-31.13:00:00 -0700"
set security authentication-key-chains key-chain secret123 key 3 secret "$9$mTQn"
set security authentication-key-chains key-chain secret123 key 3 start-time
  "2014-5-31.13:00:00 -0700"
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection minimum-interval 100
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication key-chain
  secret123
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication algorithm
  meticulous-keyed-md5

```

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

To configure IS-IS BFD authentication:

1. Configure the authentication keychain.


```

[edit security authentication-key-chains key-chain secret123]
user@R1# set description for-isis-bfd
user@R1# set key 1 secret "$9$cW-yrv"
user@R1# set key 1 start-time "2012-5-31.13:00:00 -0700"
user@R1# set key 2 secret "$9$m5T3"
user@R1# set key 2 start-time "2013-5-31.13:00:00 -0700"
user@R1# set key 3 secret "$9$mTQn"
user@R1# set key 3 start-time "2014-5-31.13:00:00 -0700"

```
2. Enable BFD.


```

[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]
user@R1# set minimum-interval 100

```
3. Apply the authentication keychain.


```

[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]

```

```
user@R1# set authentication key-chain secret123
```

4. Set the authentication type.

```
[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]  
user@R1# set authentication algorithm meticulous-keyed-md5
```

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

```
user@R1# show protocols  
isis {  
  interface ge-1/2/0.0 {  
    bfd-liveness-detection {  
      minimum-interval 100;  
      authentication {  
        key-chain secret123;  
        algorithm meticulous-keyed-md5;  
      }  
    }  
  }  
}  
  
user@R1# show security  
authentication-key-chains {  
  key-chain secret123 {  
    description for-isis-bfd;  
    key 1 {  
      secret "$9$cW-yrv"; ## SECRET-DATA  
      start-time "2012-5-31.13:00:00 -0700";  
    }  
    key 2 {  
      secret "$9$m5T3"; ## SECRET-DATA  
      start-time "2013-5-31.13:00:00 -0700";  
    }  
    key 3 {  
      secret "$9$mTQn"; ## SECRET-DATA  
      start-time "2014-5-31.13:00:00 -0700";  
    }  
  }  
}
```

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

Verification

Confirm that the configuration is working properly.

Verifying IS-IS BFD Authentication

Purpose Verify the status of IS-IS BFD authentication.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@R1> show bfd session extensive
```


Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Down	ge-1/2/0.0	0.300	1.000	3

Client ISIS L1, TX interval 0.100, RX interval 0.100, Authenticate
keychain secret123, algo meticulous-keyed-md5, mode strict
 Client ISIS L2, TX interval 0.100, RX interval 0.100, Authenticate
keychain secret123, algo meticulous-keyed-md5, mode strict
 Session down time 00:35:13, previous up time 00:12:17
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Logical system 2, routing table index 85
 Min async interval 0.100, min slow interval 1.000
 Adaptive async TX interval 0.100, RX interval 0.100
 Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3
 Remote min TX interval 0.100, min RX interval 0.100, multiplier 3
 Local discriminator 2, remote discriminator 1
 Echo mode disabled/inactive, no-absorb, no-refresh
Authentication enabled/active, keychain secret123, algo meticulous-keyed-md5, mode strict
 Session ID: 0x100101

1 sessions, 2 clients
 Cumulative transmit rate 1.0 pps, cumulative receive rate 10.0 pps

Meaning The output shows that BFD authentication is enabled on IS-IS Level 1 and Level 2.

Related Documentation

- [Example: Configuring BFD for IS-IS on page 91](#)

IS-IS Multitopology Routing and IPv6 Support

- [Example: Configuring IS-IS Multicast Topology on page 109](#)
- [Example: Configuring IS-IS Dual Stacking of IPv4 and IPv6 Unicast Addresses on page 124](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 131](#)

Example: Configuring IS-IS Multicast Topology

- [IS-IS Multicast Topologies Overview on page 109](#)
- [Example: Configuring IS-IS Multicast Topology on page 110](#)

IS-IS Multicast Topologies Overview

Most multicast routing protocols perform a reverse-path forwarding (RPF) check on the source of multicast data packets. If a packet comes in on the interface that is used to send data to the source, the packet is accepted and forwarded to one or more downstream interfaces. Otherwise, the packet is discarded and a notification is sent to the multicast routing protocol running on the interface.

In certain instances, the unicast routing table used for the RPF check is also the table used for forwarding unicast data packets. Thus, unicast and multicast routing are congruent. In other cases, where it is preferred that multicast routing be independent of unicast routing, the multicast routing protocols are configured to perform the RPF check using an alternate unicast routing table `inet.2`.

You can configure IS-IS to calculate an alternate IPv4 multicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to `inet.2`. The IS-IS interface metrics for the multicast topology can be configured independently of the unicast metrics. You can also selectively disable interfaces from participating in the multicast topology while continuing to participate in the regular unicast topology. This enables you to exercise control over the paths that multicast data takes through a network so that it is independent of unicast data paths. You can also configure IS-IS to calculate an alternate IPv6 multicast topology, in addition to the normal IPv6 unicast topology.



NOTE: IS-IS only starts advertising the routes when the interface routes are in `inet.2`.

[Table 4 on page 110](#) lists the various IPv4 statements you can use to configure IS-IS topologies.

Table 4: IPv4 Statements

Statement	Description
<code>ipv4-multicast</code>	Enables an alternate IPv4 multicast topology.
<code>ipv4-multicast-metric <i>number</i></code>	Configures the multicast metric for an alternate IPv4 multicast topology.
<code>no-ipv4-multicast</code>	Excludes an interface from the IPv4 multicast topology.
<code>no-unicast-topology</code>	Excludes an interface from the IPv4 unicast topologies.

[Table 5 on page 110](#) lists the various IPv6 statements you can use to configure IS-IS topologies.

Table 5: IPv6 Statements

Statement	Description
<code>ipv6-multicast</code>	Enables an alternate IPv6 multicast topology.
<code>ipv6-multicast-metric <i>number</i></code>	Configures the multicast metric for an alternate IPv6 multicast topology.
<code>ipv6-unicast-metric <i>number</i></code>	Configures the unicast metric for an alternate IPv6 multicast topology.
<code>no-ipv6-multicast</code>	Excludes an interface from the IPv6 multicast topology.
<code>no-ipv6-unicast</code>	Excludes an interface from the IPv6 unicast topologies.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

Example: Configuring IS-IS Multicast Topology

This example shows how to configure a multicast topology for an IS-IS network.

- [Requirements on page 110](#)
- [Overview on page 111](#)
- [Configuration on page 111](#)
- [Verification on page 115](#)

Requirements

Before you begin, configure IS-IS on all routers. See [“Example: Configuring IS-IS” on page 14](#) for information about the required IS-IS configuration.

This example uses the following hardware and software components:


- Junos OS Release 7.3 or later
- M Series, MX Series, and T Series routers

Overview

This example shows an IS-IS multicast topology configuration. Three routers are connected to each other. A loopback interface is configured on each router.

Figure 17 on page 111 shows the sample network.

Figure 17: Configuring IS-IS Multicast Topology



Configuration

CLI Quick Configuration

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

Router R1

```
set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 15
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 18
set protocols isis interface so-0/0/0 level 2 metric 20
set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 14
set protocols isis interface so-1/0/0 level 1 metric 13
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 12
```

```
set protocols isis interface so-1/0/0 level 2 metric 29
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 23
set protocols isis interface fxp0.0 disable
```

Router R2

```
set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 13
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 12
set protocols isis interface so-0/0/0 level 2 metric 29
set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 23
set protocols isis interface so-1/0/0 level 1 metric 14
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 18
set protocols isis interface so-1/0/0 level 2 metric 32
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 26
set protocols isis interface fxp0.0 disable
```

Router R3

```
set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 19
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 11
set protocols isis interface so-0/0/0 level 2 metric 27
set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 21
set protocols isis interface so-1/0/0 level 1 metric 16
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 26
set protocols isis interface so-1/0/0 level 2 metric 30
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 20
set protocols isis interface fxp0.0 disable
```

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

To configure IS-IS multicast topologies:

1. Enable the multicast topology for IS-IS by using the **ipv4-multicast** statement.

Routers R1, R2, and R3

```
[edit protocols isis]
user@host# set traceoptions file isis size 5m world-readable
user@host# set traceoptions flag error
user@host# set topologies ipv4-multicast
```

2. Enable multicast metrics on the first SONET/SDH Interface by using the **ipv4-multicast-metric** statement.

Router R1

```
[edit protocols isis interface so-0/0/0 ]
user@R1# set level 1 metric 15
user@R1# set level 1 ipv4-multicast-metric 18
user@R1# set level 2 metric 20
```

```
user@R1# set level 2 ipv4-multicast-metric 14
```

Router R2

```
[edit protocols isis interface so-0/0/0]
user@R2# set level 1 metric 13
user@R2# set level 1 ipv4-multicast-metric 12
user@R2# set level 2 metric 29
user@R2# set level 2 ipv4-multicast-metric 23
```

Router R3

```
[edit protocols isis interface so-0/0/0]
user@R3# set level 1 metric 19
user@R3# set level 1 ipv4-multicast-metric 11
user@R3# set level 2 metric 27
user@R3# set level 2 ipv4-multicast-metric 21
```

3. Enable multicast metrics on a second sonet Interface by using the **ipv4-multicast-metric** statement.

Router R1

```
[edit protocols isis interface so-1/0/0]
user@R1# set level 1 metric 13
user@R1# set level 1 ipv4-multicast-metric 12
user@R1# set level 2 metric 29
user@R1# set level 2 ipv4-multicast-metric 23
```

Router R2

```
[edit protocols isis interface so-1/0/0]
user@R2# set level 1 metric 14
user@R2# set level 1 ipv4-multicast-metric 18
user@R2# set level 2 metric 32
user@R2# set level 2 ipv4-multicast-metric 26
```

Router R3

```
[edit protocols isis interface so-1/0/0]
user@R3# set level 1 metric 16
user@R3# set level 1 ipv4-multicast-metric 26
user@R3# set level 2 metric 30
user@R3# set level 2 ipv4-multicast-metric 20
```

4. Disable the out-of-band management port, fxp0.

Routers R1, R2, and R3

```
[edit protocols isis]
user@host# set interface fxp0.0 disable
```

5. If you are done configuring the routers, commit the configuration.

Routers R1, R2, and R3

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

Results From configuration mode, confirm your configuration by using the **show protocols isis** statement. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Router R1

```
user@R1# show protocols isis
```

```
traceoptions {
  file isis size 5m world-readable;
  flag error;
}
topologies ipv4-multicast;
interface so-0/0/0 {
  level 1 {
    metric 15;
    ipv4-multicast-metric 18;
  }
  level 2 {
    metric 20;
    ipv4-multicast-metric 14;
  }
}
interface so-1/0/0 {
  level 1 {
    metric 13;
    ipv4-multicast-metric 12;
  }
  level 2 {
    metric 29;
    ipv4-multicast-metric 23;
  }
}
interface fxp0.0 {
  disable;
}
```

Router R2

```
user@R2# show protocols isis
```

```
traceoptions {
  file isis size 5m world-readable;
  flag error;
}
topologies ipv4-multicast;
interface so-0/0/0 {
  level 1 {
    metric 13;
    ipv4-multicast-metric 12;
  }
  level 2 {
    metric 29;
    ipv4-multicast-metric 23;
  }
}
interface so-1/0/0 {
  level 1 {
    metric 14;
    ipv4-multicast-metric 18;
  }
}
```



```

        level 2 {
            metric 32;
            ipv4-multicast-metric 26;
        }
    }
    interface fxp0.0 {
        disable;
    }
}

```

Router R3

user@R3# show protocols isis

```

traceoptions {
    file isis size 5m world-readable;
    flag error;
}
topologies ipv4-multicast;
interface so-0/0/0 {
    level 1 {
        metric 19;
        ipv4-multicast-metric 11;
    }
    level 2 {
        metric 27;
        ipv4-multicast-metric 21;
    }
}
interface so-1/0/0 {
    level 1 {
        metric 16;
        ipv4-multicast-metric 26;
    }
    level 2 {
        metric 30;
        ipv4-multicast-metric 20;
    }
}
interface fxp0.0 {
    disable;
}

```

Verification

Confirm that the configuration is working properly.

- [Verifying the Connection Between Routers R1, R2, and R3 on page 115](#)
- [Verifying That IS-IS Is Configured on page 117](#)
- [Verifying the Configured Multicast Metric Values on page 119](#)
- [Verifying the Configuration of the Multicast Topology on page 120](#)

Verifying the Connection Between Routers R1, R2, and R3

Purpose Make sure that Routers R1, R2, and R3 are connected to each other.

Action Ping the other two routers from any router, to check the connectivity between the three routers as per the network topology.

```
user@R1> ping 10.0.3.9
```

```
PING 10.0.3.9 (10.0.3.9): 56 data bytes
64 bytes from 10.0.3.9: icmp_seq=0 ttl=64 time=1.299 ms
64 bytes from 10.0.3.9: icmp_seq=1 ttl=64 time=52.304 ms
64 bytes from 10.0.3.9: icmp_seq=2 ttl=64 time=1.271 ms
64 bytes from 10.0.3.9: icmp_seq=3 ttl=64 time=1.343 ms
64 bytes from 10.0.3.9: icmp_seq=4 ttl=64 time=1.434 ms
64 bytes from 10.0.3.9: icmp_seq=5 ttl=64 time=1.306 ms
^C
--- 10.0.3.9 ping statistics ---
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.271/9.826/52.304/18.997 ms
```

```
user@R1> ping 10.0.3.10
```

```
PING 10.0.3.10 (10.0.3.10): 56 data bytes
64 bytes from 10.0.3.10: icmp_seq=0 ttl=64 time=1.431 ms
64 bytes from 10.0.3.10: icmp_seq=1 ttl=64 time=1.296 ms
64 bytes from 10.0.3.10: icmp_seq=2 ttl=64 time=1.887 ms
^C
--- 10.0.3.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.296/1.538/1.887/0.253 ms
```

```
user@R2> ping 10.0.2.9
```

```
PING 10.0.2.9 (10.0.2.9): 56 data bytes
64 bytes from 10.0.2.9: icmp_seq=0 ttl=64 time=1.365 ms
64 bytes from 10.0.2.9: icmp_seq=1 ttl=64 time=1.813 ms
64 bytes from 10.0.2.9: icmp_seq=2 ttl=64 time=1.290 ms
^C
--- 10.0.2.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.290/1.489/1.813/0.231 ms
```

```
user@R2> ping 10.0.2.10
```

```
PING 10.0.2.10 (10.0.2.10): 56 data bytes
64 bytes from 10.0.2.10: icmp_seq=0 ttl=63 time=1.318 ms
64 bytes from 10.0.2.10: icmp_seq=1 ttl=63 time=1.394 ms
64 bytes from 10.0.2.10: icmp_seq=2 ttl=63 time=1.366 ms
64 bytes from 10.0.2.10: icmp_seq=3 ttl=63 time=1.305 ms
^C
--- 10.0.2.10 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.305/1.346/1.394/0.036 ms
```

```
user@R3> ping 10.0.1.10
```

```
PING 10.0.1.10 (10.0.1.10): 56 data bytes
64 bytes from 10.0.1.10: icmp_seq=0 ttl=63 time=1.316 ms
64 bytes from 10.0.1.10: icmp_seq=1 ttl=63 time=1.418 ms
64 bytes from 10.0.1.10: icmp_seq=2 ttl=63 time=1.277 ms
^C
--- 10.0.1.10 ping statistics ---
```

```
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.277/1.337/1.418/0.059 ms
```

```
user@R3> ping 10.0.1.9
```

```
PING 10.0.1.9 (10.0.1.9): 56 data bytes
64 bytes from 10.0.1.9: icmp_seq=0 ttl=64 time=1.381 ms
64 bytes from 10.0.1.9: icmp_seq=1 ttl=64 time=1.499 ms
64 bytes from 10.0.1.9: icmp_seq=2 ttl=64 time=1.300 ms
64 bytes from 10.0.1.9: icmp_seq=3 ttl=64 time=1.397 ms
^C
--- 10.0.1.9 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.300/1.394/1.499/0.071 ms
```

Meaning Routers R1, R2, and R3 have a peer relationship with each other.

Verifying That IS-IS Is Configured

Purpose Make sure that the IS-IS instance is running on Routers R1, R2, and R3, and that they are adjacent to each other.

Action Use the `show isis adjacency detail` command to check the adjacency between the routers.

Router R1

```
user@R1> show isis adjacency detail
```

R2

```
Interface: so-0/0/0, Level: 1, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:23:59 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.10
```

R2

```
Interface: so-0/0/0, Level: 2, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:23:58 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.10
```

R3

```
Interface: so-1/0/0, Level: 1, State: Up, Expires in 7 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:24:20 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.02, IP addresses: 10.0.2.10
```

R3

```
Interface: so-1/0/0, Level: 2, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:24:20 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.02, IP addresses: 10.0.2.10
```

Router R2

user@R2> show isis adjacency detail

R1

Interface: so-0/0/0, Level: 1, State: Up, Expires in 20 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:50 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.9

R1

Interface: so-0/0/0, Level: 2, State: Up, Expires in 26 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:50 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.9

R3

Interface: so-1/0/0, Level: 1, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:22 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.03, IP addresses: 10.0.3.10

R3

Interface: so-1/0/0, Level: 2, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:22 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.03, IP addresses: 10.0.3.10

Router R3

user@R3> show isis adjacency detail

R2

Interface: so-0/0/0, Level: 1, State: Up, Expires in 18 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:09 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.03, IP addresses: 10.0.3.9

R2

Interface: so-0/0/0, Level: 2, State: Up, Expires in 22 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:09 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.03, IP addresses: 10.0.3.9

R1

Interface: so-1/0/0, Level: 1, State: Up, Expires in 21 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:59 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast

Restart capable: Yes, Adjacency advertisement: Advertise
 LAN id: R3.02, IP addresses: 10.0.2.9

R1

Interface: so-1/0/0, Level: 2, State: Up, Expires in 19 secs
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:59 ago
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
 Topologies: IPV4-Multicast
 Restart capable: Yes, Adjacency advertisement: Advertise
 LAN id: R3.02, IP addresses: 10.0.2.9

Meaning IS-IS is configured on Routers R1, R2, and R3, and they are adjacent to each other.

Verifying the Configured Multicast Metric Values

Purpose Make sure that the SPF calculations are accurate as per the configured multicast metric values on Routers R1, R2, and R3.

Action Use the `show isis spf results` command to check the SPF calculations for the network.

Router R1

user@R1> show isis spf results

```
...
IPv4 Multicast IS-IS level 1 SPF results:
Node Metric Interface NH Via SNPA
R3.03 28 so-1/0/0 IPV4 R3 0:1b:c0:86:54:bd
R2.00 18 so-0/0/0 IPV4 R2 0:1b:c0:86:54:bd
R3.00 17 so-1/0/0 IPV4 R3 0:1b:c0:86:54:bd
R1.00 0
4 nodes
```

```
IPv4 Multicast IS-IS level 2 SPF results:
Node Metric Interface NH Via SNPA
R3.03 40 so-0/0/0 IPV4 R2 0:1b:c0:86:54:bd
R3.00 22 so-1/0/0 IPV4 R3 0:1b:c0:86:54:bd
R2.00 14 so-0/0/0 IPV4 R2 0:1b:c0:86:54:bd
R1.00 0
4 nodes
```

Router R2

user@R2> show isis spf results

```
...
IPv4 Multicast IS-IS level 1 SPF results:
Node Metric Interface NH Via SNPA
R3.02 29 so-0/0/0 IPV4 R1 0:1b:c0:86:54:bc
R3.00 18 so-1/0/0 IPV4 R3 0:1b:c0:86:54:bd
R1.00 12 so-0/0/0 IPV4 R1 0:1b:c0:86:54:bc
R2.02 12
R2.00 0
5 nodes
```

```
IPv4 Multicast IS-IS level 2 SPF results:
Node Metric Interface NH Via SNPA
R3.02 45 so-0/0/0 IPV4 R1 0:1b:c0:86:54:bc
R3.00 26 so-1/0/0 IPV4 R3 0:1b:c0:86:54:bd
R1.00 23 so-0/0/0 IPV4 R1 0:1b:c0:86:54:bc
R2.02 23
```

```
R2.00 0
      5 nodes
```

Router R3

```
user@R3> show isis spf results
```

```
...
IPv4 Multicast IS-IS level 1 SPF results:
Node Metric Interface NH Via SNPA
R3.02 26
R1.00 23 so-0/0/0 IPv4 R2 0:1b:c0:86:54:bc
R2.02 23 so-0/0/0 IPv4 R2 0:1b:c0:86:54:bc
R2.00 11 so-0/0/0 IPv4 R2 0:1b:c0:86:54:bc
R3.03 11
R3.00 0
      6 nodes
```

```
IPv4 Multicast IS-IS level 2 SPF results:
Node Metric Interface NH Via SNPA
R2.02 34 so-1/0/0 IPv4 R1 0:1b:c0:86:54:bc
R2.00 21 so-0/0/0 IPv4 R2 0:1b:c0:86:54:bc
R3.03 21
R1.00 20 so-1/0/0 IPv4 R1 0:1b:c0:86:54:bc
R3.02 20
R3.00 0
      6 nodes
```

Meaning The configured multicast metric values are used in SPF calculations for the IS-IS network.

Verifying the Configuration of the Multicast Topology

Purpose Make sure that the multicast topology is configured on Routers R1, R2, and R3.

Action Use the **show isis database detail** command to verify the multicast topology configuration on the routers.

Router R1

```
user@R1> show isis database detail
```

```
IS-IS level 1 link-state database:
```

```
R1.00-00 Sequence: 0x142, Checksum: 0xd07, Lifetime: 663 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 15
  IPv4 Unicast IS neighbor: R3.02 Metric: 15
  IPv4 Multicast IS neighbor: R2.02 Metric: 18
  IPv4 Multicast IS neighbor: R3.02 Metric: 17
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 15 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 15 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 883 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 13
  IPv4 Unicast IS neighbor: R3.03 Metric: 14
  IPv4 Multicast IS neighbor: R2.02 Metric: 12
  IPv4 Multicast IS neighbor: R3.03 Metric: 18
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 13 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 14 Internal Up
```

```
R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 913 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
```

```

IPv4 Unicast IS neighbor: R2.00    Metric:      0

R3.00-00 Sequence: 0x13c, Checksum: 0xc8de, Lifetime: 488 secs
IPv4 Unicast IS neighbor: R3.02    Metric:      16
IPv4 Unicast IS neighbor: R3.03    Metric:      19
IPv4 Multicast IS neighbor: R3.02   Metric:      26
IPv4 Multicast IS neighbor: R3.03   Metric:      11
IP IPv4 Unicast prefix: 10.0.2.8/30 Metric:      16 Internal Up
IP IPv4 Unicast prefix: 10.0.3.8/30 Metric:      19 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 625 secs
IPv4 Unicast IS neighbor: R1.00    Metric:      0
IPv4 Unicast IS neighbor: R3.00    Metric:      0

R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 714 secs
IPv4 Unicast IS neighbor: R2.00    Metric:      0
IPv4 Unicast IS neighbor: R3.00    Metric:      0

IS-IS level 2 link-state database:

R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 816 secs
IPv4 Unicast IS neighbor: R2.02    Metric:      20
IPv4 Unicast IS neighbor: R3.02    Metric:      31
IPv4 Multicast IS neighbor: R2.02   Metric:      14
IPv4 Multicast IS neighbor: R3.02   Metric:      22
IP IPv4 Unicast prefix: 10.0.1.8/30 Metric:      20 Internal Up
IP IPv4 Unicast prefix: 10.0.2.8/30 Metric:      31 Internal Up
IP IPv4 Unicast prefix: 10.0.3.8/30 Metric:      29 Internal Up

R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 966 secs
IPv4 Unicast IS neighbor: R2.02    Metric:      29
IPv4 Unicast IS neighbor: R3.03    Metric:      32
IPv4 Multicast IS neighbor: R2.02   Metric:      23
IPv4 Multicast IS neighbor: R3.03   Metric:      26
IP IPv4 Unicast prefix: 10.0.1.8/30 Metric:      29 Internal Up
IP IPv4 Unicast prefix: 10.0.2.8/30 Metric:      28 Internal Up
IP IPv4 Unicast prefix: 10.0.3.8/30 Metric:      32 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 966 secs
IPv4 Unicast IS neighbor: R1.00    Metric:      0
IPv4 Unicast IS neighbor: R2.00    Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 805 secs
IPv4 Unicast IS neighbor: R3.02    Metric:      30
IPv4 Unicast IS neighbor: R3.03    Metric:      27
IPv4 Multicast IS neighbor: R3.02   Metric:      20
IPv4 Multicast IS neighbor: R3.03   Metric:      21
IP IPv4 Unicast prefix: 10.0.1.8/30 Metric:      31 Internal Up
IP IPv4 Unicast prefix: 10.0.2.8/30 Metric:      30 Internal Up
IP IPv4 Unicast prefix: 10.0.3.8/30 Metric:      27 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 844 secs
IPv4 Unicast IS neighbor: R1.00    Metric:      0
IPv4 Unicast IS neighbor: R3.00    Metric:      0

R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 844 secs
IPv4 Unicast IS neighbor: R2.00    Metric:      0
IPv4 Unicast IS neighbor: R3.00    Metric:      0

```

Router R2

```
user@R2> show isis database detail
```

```
IS-IS level 1 link-state database:
```

```
R1.00-00 Sequence: 0x142, Checksum: 0xd07, Lifetime: 524 secs
  IPv4 Unicast IS neighbor: R2.02      Metric:      15
  IPv4 Unicast IS neighbor: R3.02      Metric:      15
  IPv4 Multicast IS neighbor: R2.02    Metric:      18
  IPv4 Multicast IS neighbor: R3.02    Metric:      17
  IP IPv4 Unicast prefix: 10.0.1.8/30  Metric:      15 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30  Metric:      15 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 748 secs
  IPv4 Unicast IS neighbor: R2.02      Metric:      13
  IPv4 Unicast IS neighbor: R3.03      Metric:      14
  IPv4 Multicast IS neighbor: R2.02    Metric:      12
  IPv4 Multicast IS neighbor: R3.03    Metric:      18
  IP IPv4 Unicast prefix: 10.0.1.8/30  Metric:      13 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30  Metric:      14 Internal Up
```

```
R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 777 secs
  IPv4 Unicast IS neighbor: R1.00      Metric:       0
  IPv4 Unicast IS neighbor: R2.00      Metric:       0
```

```
R3.00-00 Sequence: 0x13d, Checksum: 0xc6df, Lifetime: 1102 secs
  IPv4 Unicast IS neighbor: R3.02      Metric:      16
  IPv4 Unicast IS neighbor: R3.03      Metric:      19
  IPv4 Multicast IS neighbor: R3.02    Metric:      26
  IPv4 Multicast IS neighbor: R3.03    Metric:      11
  IP IPv4 Unicast prefix: 10.0.2.8/30  Metric:      16 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30  Metric:      19 Internal Up
```

```
R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 488 secs
  IPv4 Unicast IS neighbor: R1.00      Metric:       0
  IPv4 Unicast IS neighbor: R3.00      Metric:       0
```

```
R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 577 secs
  IPv4 Unicast IS neighbor: R2.00      Metric:       0
  IPv4 Unicast IS neighbor: R3.00      Metric:       0
```

```
IS-IS level 2 link-state database:
```

```
R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 676 secs
  IPv4 Unicast IS neighbor: R2.02      Metric:      20
  IPv4 Unicast IS neighbor: R3.02      Metric:      31
  IPv4 Multicast IS neighbor: R2.02    Metric:      14
  IPv4 Multicast IS neighbor: R3.02    Metric:      22
  IP IPv4 Unicast prefix: 10.0.1.8/30  Metric:      20 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30  Metric:      31 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30  Metric:      29 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 831 secs
  IPv4 Unicast IS neighbor: R2.02      Metric:      29
  IPv4 Unicast IS neighbor: R3.03      Metric:      32
  IPv4 Multicast IS neighbor: R2.02    Metric:      23
  IPv4 Multicast IS neighbor: R3.03    Metric:      26
  IP IPv4 Unicast prefix: 10.0.1.8/30  Metric:      29 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30  Metric:      28 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30  Metric:      32 Internal Up
```



```

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 831 secs
  IPv4 Unicast IS neighbor: R1.00      Metric:      0
  IPv4 Unicast IS neighbor: R2.00      Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 667 secs
  IPv4 Unicast IS neighbor: R3.02      Metric:      30
  IPv4 Unicast IS neighbor: R3.03      Metric:      27
  IPv4 Multicast IS neighbor: R3.02     Metric:      20
  IPv4 Multicast IS neighbor: R3.03     Metric:      21
  IP IPv4 Unicast prefix: 10.0.1.8/30  Metric:      31 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30  Metric:      30 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30  Metric:      27 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 707 secs
  IPv4 Unicast IS neighbor: R1.00      Metric:      0
  IPv4 Unicast IS neighbor: R3.00      Metric:      0

R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 707 secs
  IPv4 Unicast IS neighbor: R2.00      Metric:      0
  IPv4 Unicast IS neighbor: R3.00      Metric:      0

```

Router R3

```
user@R3> show isis database detail
```

```
IS-IS level 1 link-state database:
```

```

R1.00-00 Sequence: 0x143, Checksum: 0xb08, Lifetime: 1155 secs
  IPv4 Unicast IS neighbor: R2.02      Metric:      15
  IPv4 Unicast IS neighbor: R3.02      Metric:      15
  IPv4 Multicast IS neighbor: R2.02     Metric:      18
  IPv4 Multicast IS neighbor: R3.02     Metric:      17
  IP IPv4 Unicast prefix: 10.0.1.8/30  Metric:      15 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30  Metric:      15 Internal Up

R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 687 secs
  IPv4 Unicast IS neighbor: R2.02      Metric:      13
  IPv4 Unicast IS neighbor: R3.03      Metric:      14
  IPv4 Multicast IS neighbor: R2.02     Metric:      12
  IPv4 Multicast IS neighbor: R3.03     Metric:      18
  IP IPv4 Unicast prefix: 10.0.1.8/30  Metric:      13 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30  Metric:      14 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 716 secs
  IPv4 Unicast IS neighbor: R1.00      Metric:      0
  IPv4 Unicast IS neighbor: R2.00      Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0xc6df, Lifetime: 1044 secs
  IPv4 Unicast IS neighbor: R3.02      Metric:      16
  IPv4 Unicast IS neighbor: R3.03      Metric:      19
  IPv4 Multicast IS neighbor: R3.02     Metric:      26
  IPv4 Multicast IS neighbor: R3.03     Metric:      11
  IP IPv4 Unicast prefix: 10.0.2.8/30  Metric:      16 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30  Metric:      19 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 430 secs
  IPv4 Unicast IS neighbor: R1.00      Metric:      0
  IPv4 Unicast IS neighbor: R3.00      Metric:      0

R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 519 secs
  IPv4 Unicast IS neighbor: R2.00      Metric:      0

```

```
IPv4 Unicast IS neighbor: R3.00    Metric:      0
```

IS-IS level 2 link-state database:

```
R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 617 secs
```

```
IPv4 Unicast IS neighbor: R2.02    Metric:      20
IPv4 Unicast IS neighbor: R3.02    Metric:      31
IPv4 Multicast IS neighbor: R2.02   Metric:      14
IPv4 Multicast IS neighbor: R3.02   Metric:      22
IP IPv4 Unicast prefix: 10.0.1.8/30 Metric:      20 Internal Up
IP IPv4 Unicast prefix: 10.0.2.8/30 Metric:      31 Internal Up
IP IPv4 Unicast prefix: 10.0.3.8/30 Metric:      29 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 769 secs
```

```
IPv4 Unicast IS neighbor: R2.02    Metric:      29
IPv4 Unicast IS neighbor: R3.03    Metric:      32
IPv4 Multicast IS neighbor: R2.02   Metric:      23
IPv4 Multicast IS neighbor: R3.03   Metric:      26
IP IPv4 Unicast prefix: 10.0.1.8/30 Metric:      29 Internal Up
IP IPv4 Unicast prefix: 10.0.2.8/30 Metric:      28 Internal Up
IP IPv4 Unicast prefix: 10.0.3.8/30 Metric:      32 Internal Up
```

```
R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 769 secs
```

```
IPv4 Unicast IS neighbor: R1.00    Metric:      0
IPv4 Unicast IS neighbor: R2.00    Metric:      0
```

```
R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 610 secs
```

```
IPv4 Unicast IS neighbor: R3.02    Metric:      30
IPv4 Unicast IS neighbor: R3.03    Metric:      27
IPv4 Multicast IS neighbor: R3.02   Metric:      20
IPv4 Multicast IS neighbor: R3.03   Metric:      21
IP IPv4 Unicast prefix: 10.0.1.8/30 Metric:      31 Internal Up
IP IPv4 Unicast prefix: 10.0.2.8/30 Metric:      30 Internal Up
IP IPv4 Unicast prefix: 10.0.3.8/30 Metric:      27 Internal Up
```

```
R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 649 secs
```

```
IPv4 Unicast IS neighbor: R1.00    Metric:      0
IPv4 Unicast IS neighbor: R3.00    Metric:      0
```

```
R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 649 secs
```

```
IPv4 Unicast IS neighbor: R2.00    Metric:      0
IPv4 Unicast IS neighbor: R3.00    Metric:      0
```

Meaning Multicast topology is configured on Routers R1, R2, and R3.

Related Documentation

- [Example: Configuring Multitopology Routing Based on a Multicast Source](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 131](#)

Example: Configuring IS-IS Dual Stacking of IPv4 and IPv6 Unicast Addresses

- [Understanding Dual Stacking on page 125](#)
- [Example: Configuring IS-IS Dual Stacking of IPv4 and IPv6 Unicast Addresses on page 126](#)

Understanding Dual Stacking

Service providers and enterprises are faced with growing their networks using IPv6, while continuing to serve IPv4 customers.

Increasingly, the public side of network address translation (NAT) devices is IPv6 rather than IPv4. Service providers cannot continue giving customers globally routable IPv4 addresses, they cannot get new globally routable IPv4 addresses for expanding their own networks, and yet they must continue to serve both IPv4 customers and new customers, all of whom are primarily trying to reach IPv4 destinations.

IPv4 and IPv6 must coexist for some number of years, and their coexistence must be transparent to end users. If an IPv4-to-IPv6 transition is successful, end users should not even notice it.

A dual-stack device is a device with network interfaces that can originate and understand both IPv4 and IPv6 packets.

Other strategies, such as manually or dynamically configured tunnels and translation devices exist, but dual stacking is often the preferable solution in many scenarios. The dual-stacked device can interoperate equally with IPv4 devices, IPv6 devices, and other dual-stacked devices. When both devices are dual stacked, the two devices agree on which IP version to use.

The transition is driven by DNS. If a dual-stacked device queries the name of a destination and DNS gives it an IPv4 address (a DNS A Record), it sends IPv4 packets. If DNS responds with an IPv6 address (a DNS AAAA Record), it sends IPv6 packets.

Keep in mind that if you are going to dual stack all of your network devices, the interfaces need both an IPv6 and an IPv4 address. This raises the issue that the Internet has run out of IPv4 addresses, which is the main reason IPv6 is needed in the first place. If you do not have an abundant supply of IPv4 addresses to apply to your devices, you can still use dual stacking, but you will need to conserve your supply of IPv4 addresses by using network address translation (NAT). Building dual-stacked networks with a mix of global IPv6 addresses and NAT-ed IPv4 addresses is quite feasible. Some specific solutions include carrier-grade NAT (CGN), NAT44(4), NAT64, NAT464, and dual-stack lite.

[Table 6 on page 125](#) describes at a high level how to pick a network addressing technique. In reality, a complete solution might include a set of techniques to satisfy multiple service needs. It is important to understand the backbone technology being used on the network and also to know if the provider has control over the access customer premises equipment (CPE).

Table 6: Choosing the Right Solution to Address Next-Generation Addressing Requirements

CPE Network	Access Network	Destinations	Solution
IPv4	IPv4	IPv4 Internet	NAT44(4)
IPv4/IPv6	IPv6	IPv4 Internet	DS-Lite with NAT44

Table 6: Choosing the Right Solution to Address Next-Generation Addressing Requirements (*continued*)

CPE Network	Access Network	Destinations	Solution
IPv4/IPv6	IPv4	IPv6 Internet	6rd (6to4)
IPv4	IPv6	IPv4 Internet	NAT64

Example: Configuring IS-IS Dual Stacking of IPv4 and IPv6 Unicast Addresses

This example shows how to configure IPv4 and IPv6 dual stacking in IS-IS.

- [Requirements on page 126](#)
- [Overview on page 126](#)
- [Configuration on page 127](#)
- [Verification on page 129](#)

Requirements

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

Overview**Video: IS-IS Dual Stacking**

You can use IPv4 and IPv6 dual stacking to begin your migration from IPv4 to IPv6 by implementing IPv6 alongside IPv4 in your existing networks. This allows you to implement IPv6 so that you can provide the same services over IPv6—for example, video, voice, high-quality data—that you currently provide in your IPv4 networks. You can then perform incremental upgrades to IPv6 and avoid service disruptions while migrating from IPv4 to IPv6.

Unlike RIP and OSPF, IS-IS does not require a distinct protocol or a new version to support IPv6. Because IS-IS uses ISO addresses, the configuration for IPv6 and IPv4 is identical in the Junos OS implementation of IS-IS. For IS-IS to carry IPv6 routes, you only need to add IPv6 addresses to IS-IS enabled interfaces or include other IPv6 routes in your IS-IS export policy.

The only explicit configuration needed in IS-IS with regard to IPv6 is if you want to disable it. Alternatively, you can disable IPv4 routing and use IS-IS with IPv6 only. An example of each is provided here:

Disable IPv6 routing in IS-IS:


```
[edit protocols isis]
user@host# set no-ipv6-routing
```

Use IS-IS exclusively for IPv6 routing:

```
[edit protocols isis]
user@host# set no-ipv4-routing
```

Figure 18 on page 127 shows the topology used in this example.

Figure 18: IS-IS IPv4 and IPv6 Dual Stacking Topology



“CLI Quick Configuration” on page 127 shows the configuration for all of the devices in Figure 18 on page 127. The section “Step-by-Step Procedure” on page 128 describes the steps on Device R1.

Configuration

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

Device R1

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::1/128
set protocols isis interface fe-1/2/0.0
set protocols isis interface fe-1/2/1.0
set protocols isis interface lo0.0
```

Device R2

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::2/128
set protocols isis interface fe-1/2/0.0
set protocols isis interface lo0.0
```

```

Device R3    set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.18/30
              set interfaces fe-1/2/0 unit 0 family iso
              set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
              set interfaces lo0 unit 0 family inet address 192.168.0.3/32
              set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0003.00
              set interfaces lo0 unit 0 family inet6 address 2001:db8::3/128
              set protocols isis interface fe-1/2/0.0
              set protocols isis interface lo0.0

```

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

To configure IS-IS dual stacking:

1. Configure the interfaces, including both IPv4 and IPv6 addresses on each interface.

Optionally, include the **eui-64** statement to automatically generate the host number portion of interface addresses.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 0 family inet address 10.0.0.1/30
user@R1# set fe-1/2/0 unit 0 family iso
user@R1# set fe-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
user@R1# set fe-1/2/1 unit 0 family inet address 10.0.0.17/30
user@R1# set fe-1/2/1 unit 0 family iso
user@R1# set fe-1/2/1 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1# set lo0 unit 0 family inet address 192.168.0.1/32
user@R1# set lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
user@R1# set lo0 unit 0 family inet6 address 2001:db8::1/128

```

2. Enable IS-IS on the interfaces.

```

[edit protocols isis]
user@R1# set interface fe-1/2/0.0
user@R1# set interface fe-1/2/1.0
user@R1# set interface lo0.0

```

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

```

user@R1# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.1/30;
    }
    family iso;
    family inet6 {
      address 2001:db8:0:5::/64 {
        eui-64;
      }
    }
  }
}

```

```

fe-1/2/1 {
  unit 0 {
    family inet {
      address 10.0.0.17/30;
    }
    family iso;
    family inet6 {
      address 2001:db8:0:1::/64 {
        eui-64;
      }
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.1/32;
    }
    family iso {
      address 49.0002.0192.0168.0001.00;
    }
    family inet6 {
      address 2001:db8::1/128;
    }
  }
}

user@R1# show protocols
isis {
  interface fe-1/2/0.0;
  interface fe-1/2/1.0;
  interface lo0.0;
}

```

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

Verification

Confirm that the configuration is working properly.

- [Checking the Neighbor Adjacencies on page 129](#)
- [Pinging the IPv6 Interfaces on page 130](#)
- [Checking the IPv6 Routing Table on page 131](#)

Checking the Neighbor Adjacencies

Purpose Determine what topologies are supported on neighboring IS-IS devices.

Action From operational mode, enter the **show isis adjacency detail** command.

```
user@R1> show isis adjacency detail
```

R2

```

Interface: fe-1/2/0.0, Level: 3, State: Up, Expires in 24 secs
Priority: 0, Up/Down transitions: 1, Last transition: 18:34:08 ago
Circuit type: 3, Speaks: IP, IPv6

```

```

Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 10.0.0.2
IPv6 addresses: fe80::2a0:a514:0:24c

```

R3

```

Interface: fe-1/2/1.0, Level: 3, State: Up, Expires in 21 secs
Priority: 0, Up/Down transitions: 1, Last transition: 18:33:41 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 10.0.0.18
IPv6 addresses: fe80::2a0:a514:0:124c

```

Meaning As expected, the output shows that the two neighbors support both IPv4 and IPv6. The IPv4 address and the IPv6 link-local address are also shown.

Pinging the IPv6 Interfaces

Purpose Make sure that you can ping the remote IPv6 interfaces.

Action From operational mode, enter the **ping** command to ping from Device R2 to Device R3.

1. Determine the IPv6 address assigned to Device R3.

If you use EUI-64 addressing as shown in the example, the host portion of the IPv6 addresses is assigned automatically. To determine what addresses are assigned, use the **show interfaces terse** command on Device R3.

```

user@R3> show interfaces terse
Interface           Admin Link Proto  Local              Remote
fe-1/2/0
fe-1/2/0.0          up    up    inet   10.0.0.18/30
                   iso
                   inet6   2001:db8:0:1:2a0:a514:0:124c/64
                   fe80::2a0:a514:0:124c/64

lo0
lo0.0               up    up    inet   192.168.0.3        --> 0/0
                   iso
                   inet6   2001:db8::3
                   fe80::2a0:a50f:fc56:14c

```

The IPv6 addresses that should be pingable are 2001:db8:0:1:2a0:a514:0:124c and 2001:db8::3.

2. From Device R2, ping the Device R3 fe-1/2/0.0 IPv6 interface address and the lo0.0 IPv6 interface address.

```

user@R2> ping 2001:db8:0:1:2a0:a514:0:124c
PING6(56=40+8+8 bytes) 2001:db8:0:5:2a0:a514:0:24c -->
2001:db8:0:1:2a0:a514:0:124c
16 bytes from 2001:db8:0:1:2a0:a514:0:124c, icmp_seq=0 hlim=63 time=2.373 ms
16 bytes from 2001:db8:0:1:2a0:a514:0:124c, icmp_seq=1 hlim=63 time=1.600 ms
16 bytes from 2001:db8:0:1:2a0:a514:0:124c, icmp_seq=2 hlim=63 time=2.228 ms

user@R2> ping 2001:db8::3
PING6(56=40+8+8 bytes) 2001:db8:0:5:2a0:a514:0:24c --> 2001:db8::3
16 bytes from 2001:db8::3, icmp_seq=0 hlim=63 time=1.797 ms
16 bytes from 2001:db8::3, icmp_seq=1 hlim=63 time=1.430 ms
16 bytes from 2001:db8::3, icmp_seq=2 hlim=63 time=2.525 ms

```


Meaning This test confirms that IS-IS has learned the IPv6 routes.

Checking the IPv6 Routing Table

Purpose Verify that the expected routes are in the IPv6 routing table.

Action user@R1> `show route table inet6.0`
 inet6.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
 + = Active Route, - = Last Active, * = Both

```

2001:db8::1/128    *[Direct/0] 18:52:52
                  > via lo0.0
2001:db8::2/128    *[IS-IS/15] 01:59:52, metric 10
                  > to fe80::2a0:a514:0:24c via fe-1/2/0.0
2001:db8::3/128    *[IS-IS/15] 01:59:52, metric 10
                  > to fe80::2a0:a514:0:124c via fe-1/2/1.0
2001:db8:0:1::/64  *[Direct/0] 18:52:15
                  > via fe-1/2/1.0
2001:db8:0:1:2a0:a514:0:114c/128
                  *[Local/0] 18:52:48
                  Local via fe-1/2/1.0
2001:db8:0:5::/64  *[Direct/0] 18:52:49
                  > via fe-1/2/0.0
2001:db8:0:5:2a0:a514:0:14c/128
                  *[Local/0] 18:52:49
                  Local via fe-1/2/0.0
fe80::/64          *[Direct/0] 18:52:49
                  > via fe-1/2/0.0
                  [Direct/0] 18:52:15
                  > via fe-1/2/1.0
fe80::2a0:a50f:fc56:14c/128
                  *[Direct/0] 18:52:52
                  > via lo0.0
fe80::2a0:a514:0:14c/128
                  *[Local/0] 18:52:49
                  Local via fe-1/2/0.0
fe80::2a0:a514:0:114c/128
                  *[Local/0] 18:52:48
                  Local via fe-1/2/1.0
  
```

Meaning The output shows the IPv6 interface routes (direct and local) and the IPv6 routes learned through IS-IS.

Related Documentation

- [Example: Configuring IS-IS on page 13](#)

Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

- [Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 132](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 132](#)

Understanding IS-IS IPv4 and IPv6 Unicast Topologies

You can configure 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. 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 enables you to exercise control over the paths that unicast data takes through a network.

A topology is the set of joined nodes. IS-IS evaluates all the paths in a single topology for each IS-IS level and uses the shortest-path-first (SPF) algorithm to determine the best path among all the feasible paths. Topology discovery and SPF calculation is performed in a protocol-neutral fashion because it is done at Layer 2 of the OSI model. If you load the topology with reachability information for a certain protocol (for example, IP), the assumption is that the circuits that are supposed to provide reachability between routing devices can carry the protocol. The SPF algorithm has a per-link orientation, not a per-address family or per-protocol orientation.

Multitopology routing enables you to override this default behavior by enabling a per-address family, per-protocol SPF calculation.

The additional CPU load associated with multiple runs of the SPF algorithm is generally not an issue with the processing power available on today's routing device control planes.

The multitopology extensions alter existing type, length, and value (TLV) tuples by adding a topology ID. Each routing device in a given topology maintains its adjacencies and runs a per-topology SPF calculation.

Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

This example shows how to configure IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology.

- [Requirements on page 132](#)
- [Overview on page 132](#)
- [Configuration on page 134](#)
- [Verification on page 138](#)

Requirements

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

Overview

This example focuses on IPv4 and IPv6 unicast topologies. 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 enables you to exercise control over the paths that unicast data takes through a network.

To enable an IPv6 unicast topology for IS-IS, include the **ipv6-unicast** statement:

```
isis {  
  topologies {  
    ipv6-unicast;  
  }  
}
```

To configure a metric for the IPv6 unicast topology, include the **ipv6-unicast-metric** statement:


```
isis {  
  interface interface-name {  
    level level-number {  
      ipv6-unicast-metric number;  
    }  
  }  
}
```

To exclude an interface from the IPv6 unicast topologies for IS-IS, include the **no-ipv6-unicast** statement:

```
isis {  
  interface interface-name {  
    no-ipv6-unicast;  
  }  
}
```

[Figure 19 on page 134](#) shows the topology used in this example. The black lines indicate link membership in the IPv6 topology. The gray lines indicate membership to the IPv4 topology. Using regular TLVs, it would not be possible to build multiple topologies and run an SPF calculation based on them. The multitopology extensions describe an extension to carry the set of supported protocols in the hello packet. After activating multitopology routing support on a link, the link carries all the topologies that the underlying circuit is able to relay.

Figure 19: IS-IS IPv4 and IPv6 Unicast Topologies



“CLI Quick Configuration” on page 134 shows the configuration for all of the devices in Figure 19 on page 134. The section “Step-by-Step Procedure” on page 136 describes the steps on Device R1.

Configuration

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

```
Device R1
set interfaces so-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces so-1/2/1 unit 0 family inet address 10.0.0.13/30
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/2 unit 0 family inet address 10.0.0.17/30
set interfaces so-1/2/2 unit 0 family iso
set interfaces so-1/2/2 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::1/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0
set protocols isis interface so-1/2/1.0 no-ipv6-unicast
set protocols isis interface so-1/2/2.0
set protocols isis interface lo0.0
```

```

Device R2    set interfaces so-1/2/0 unit 0 family inet address 10.0.0.2/30
              set interfaces so-1/2/0 unit 0 family iso
              set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
              set interfaces so-1/2/1 unit 0 family inet address 10.0.0.9/30
              set interfaces so-1/2/1 unit 0 family iso
              set interfaces so-1/2/1 unit 0 family inet6 address 2001:db8:0:4::/64 eui-64
              set interfaces lo0 unit 0 family inet address 192.168.0.2/32
              set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
              set interfaces lo0 unit 0 family inet6 address 2001:db8::2/128
              set protocols isis topologies ipv6-unicast
              set protocols isis interface so-1/2/0.2
              set protocols isis interface so-1/2/1.0
              set protocols isis interface lo0.0

Device R3    set interfaces so-1/2/0 unit 0 family inet address 10.0.0.18/30
              set interfaces so-1/2/0 unit 0 family iso
              set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
              set interfaces so-1/2/1 unit 0 family inet address 10.0.0.25/30
              set interfaces so-1/2/1 unit 0 family iso
              set interfaces lo0 unit 0 family inet address 192.168.0.3/32
              set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0003.00
              set interfaces lo0 unit 0 family inet6 address 2001:db8::3/128
              set protocols isis topologies ipv6-unicast
              set protocols isis interface so-1/2/0.0
              set protocols isis interface so-1/2/1.0 no-ipv6-unicast
              set protocols isis interface lo0.0

Device R4    set interfaces so-1/2/0 unit 0 family inet address 10.0.0.5/30
              set interfaces so-1/2/0 unit 0 family iso
              set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:2::/64 eui-64
              set interfaces so-1/2/1 unit 0 family inet address 10.0.0.10/30
              set interfaces so-1/2/1 unit 0 family iso
              set interfaces so-1/2/1 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
              set interfaces so-1/2/2 unit 0 family inet address 10.0.0.21/30
              set interfaces so-1/2/2 unit 0 family iso
              set interfaces lo0 unit 0 family inet address 192.168.0.4/32
              set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0004.00
              set interfaces lo0 unit 0 family inet6 address 2001:db8::4/128
              set protocols isis topologies ipv6-unicast
              set protocols isis interface so-1/2/0.0
              set protocols isis interface so-1/2/1.0
              set protocols isis interface so-1/2/2.0 no-ipv6-unicast
              set protocols isis interface lo0.0

Device R5    set interfaces so-1/2/0 unit 0 family inet address 10.0.0.14/30
              set interfaces so-1/2/0 unit 0 family iso
              set interfaces so-1/2/1 unit 0 family inet address 10.0.0.22/30
              set interfaces so-1/2/1 unit 0 family iso
              set interfaces so-1/2/2 unit 0 family inet address 10.0.0.26/30
              set interfaces so-1/2/2 unit 0 family iso
              set interfaces so-1/2/3 unit 0 family iso
              set interfaces so-1/2/3 unit 0 family inet6 address 2001:db8:0:3::/64 eui-64
              set interfaces lo0 unit 0 family inet address 192.168.0.5/32
              set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0005.00
              set interfaces lo0 unit 0 family inet6 address 2001:db8::5/128

```

```

set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0 no-ipv6-unicast
set protocols isis interface so-1/2/1.0 no-ipv6-unicast
set protocols isis interface so-1/2/2.0 no-ipv6-unicast
set protocols isis interface so-1/2/3.0
set protocols isis interface lo0.0

```

Device R6

```

set interfaces so-1/2/0 unit 0 family inet address 10.0.0.6/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:2::/64 eui-64
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/1 unit 0 family inet6 address 2001:db8:0:3::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.6/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0006.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::6/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0
set protocols isis interface so-1/2/1.0
set protocols isis interface lo0.0

```

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

To configure an alternate IPv6 unicast topology:

1. Configure the interfaces.

```

[edit interfaces]
user@R1# set so-1/2/0 unit 0 family inet address 10.0.0.1/30
user@R1# set so-1/2/0 unit 0 family iso
user@R1# set so-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
user@R1# set so-1/2/1 unit 0 family inet address 10.0.0.13/30
user@R1# set so-1/2/1 unit 0 family iso
user@R1# set so-1/2/2 unit 0 family inet address 10.0.0.17/30
user@R1# set so-1/2/2 unit 0 family iso
user@R1# set so-1/2/2 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1# set lo0 unit 0 family inet address 192.168.0.1/32
user@R1# set lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
user@R1# set lo0 unit 0 family inet6 address 2001:db8::1/128

```

2. Enable IS-IS on the interfaces.

```

[edit protocols isis]
user@R1# set interface so-1/2/0.0
user@R1# set interface so-1/2/1.0
user@R1# set interface so-1/2/2.0
user@R1# set interface lo0.0

```

3. Enable multitopology routing on the IS-IS interfaces.

The **ipv6-unicast** statement enables multitopology IS-IS routing on all interfaces that have **family iso** and **family inet6** configured and are listed at the **[edit protocols isis interface]** hierarchy level.

```

[edit protocols isis]
user@R1# set topologies ipv6-unicast

```

4. Disable IPv6 unicast support on a given interface.

If you do not want to run multitopology IS-IS routing for IPv6 on a given interface, you can disable multitopology routing by including the **no-ipv6-unicast** statement in the IS-IS interface configuration.

```
[edit protocols isis]
user@R1# set interface so-1/2/1.0 no-ipv6-unicast
```

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

```
user@R1# show interfaces
so-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.1/30;
    }
    family iso;
    family inet6 {
      address 2001:db8:0:5::/64 {
        eui-64;
      }
    }
  }
}
so-1/2/1 {
  unit 0 {
    family inet {
      address 10.0.0.13/30;
    }
    family iso;
  }
}
so-1/2/2 {
  unit 0 {
    family inet {
      address 10.0.0.17/30;
    }
    family iso;
    family inet6 {
      address 2001:db8:0:1::/64 {
        eui-64;
      }
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.1/32;
    }
    family iso {
      address 49.0002.0192.0168.0001.00;
```

```

    }
    family inet6 {
        address 2001:db8::1/128;
    }
}

user@R1# show protocols
isis {
    topologies ipv6-unicast;
    interface so-1/2/0.0;
    interface so-1/2/1.0 {
        no-ipv6-unicast;
    }
    interface so-1/2/2.0;
    interface lo0.0;
}

```

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

Verification

Confirm that the configuration is working properly.

- [Checking the Topologies on Neighbors on page 138](#)
- [Checking the IS-IS SPF Calculations on page 139](#)
- [Checking the Tcpdump Output on page 140](#)

Checking the Topologies on Neighbors

Purpose Determine what topologies are supported on neighboring IS-IS devices.

Action From operational mode, enter the **show isis adjacency detail** command.

```
user@R1> show isis adjacency detail
```

R2

```

Interface: so-1/2/0.0, Level: 3, State: Up, Expires in 24 secs
Priority: 0, Up/Down transitions: 1, Last transition: 05:28:16 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast, IPV6-Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 10.0.0.2
IPv6 addresses: fe80::2a0:a514:0:24c

```

R5

```

Interface: so-1/2/1.0, Level: 3, State: Up, Expires in 21 secs
Priority: 0, Up/Down transitions: 1, Last transition: 05:27:47 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 10.0.0.14

```

R3

```

Interface: so-1/2/2.0, Level: 3, State: Up, Expires in 22 secs
Priority: 0, Up/Down transitions: 1, Last transition: 05:27:25 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast, IPV6-Unicast

```



```
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 10.0.0.18
IPv6 addresses: fe80::2a0:a514:0:124c
```

Meaning As expected, the adjacency with Device R5 only supports the IPv4 unicast topology, while the adjacencies with Device R2 and Device R3 support both the IPv4 and IPv6 topologies.

Checking the IS-IS SPF Calculations

Purpose Verify that separate SPF calculations are being run for IPv4 and IPv6.

Action From operational mode, enter the **show isis spf brief** command.

```
user@R1> show isis spf brief
```

IPv4 Unicast IS-IS level 1 SPF results:

Node	Metric	Interface	NH	Via	SNPA
R6.00	20	so-1/2/1.0	IPV4 R5		
R4.00	20	so-1/2/0.0	IPV4 R2		
R5.00	10	so-1/2/1.0	IPV4 R5		
R3.00	10	so-1/2/2.0	IPV4 R3		
R2.00	10	so-1/2/0.0	IPV4 R2		
R1.00	0				

6 nodes

IPv4 Unicast IS-IS level 2 SPF results:

Node	Metric	Interface	NH	Via	SNPA
R6.00	20	so-1/2/1.0	IPV4 R5		
R4.00	20	so-1/2/0.0	IPV4 R2		
R5.00	10	so-1/2/1.0	IPV4 R5		
R3.00	10	so-1/2/2.0	IPV4 R3		
R2.00	10	so-1/2/0.0	IPV4 R2		
R1.00	0				

6 nodes

IPv6 Unicast IS-IS level 1 SPF results:

Node	Metric	Interface	NH	Via	SNPA
R5.00	40	so-1/2/0.0	IPV6 R2		
R6.00	30	so-1/2/0.0	IPV6 R2		
R4.00	20	so-1/2/0.0	IPV6 R2		
R3.00	10	so-1/2/2.0	IPV6 R3		
R2.00	10	so-1/2/0.0	IPV6 R2		
R1.00	0				

6 nodes

IPv6 Unicast IS-IS level 2 SPF results:

Node	Metric	Interface	NH	Via	SNPA
R5.00	40	so-1/2/0.0	IPV6 R2		
R6.00	30	so-1/2/0.0	IPV6 R2		
R4.00	20	so-1/2/0.0	IPV6 R2		
R3.00	10	so-1/2/2.0	IPV6 R3		
R2.00	10	so-1/2/0.0	IPV6 R2		
R1.00	0				

6 nodes

Meaning As expected, SPF calculations are being performed for IPv4 and IPv6 topologies.

Checking the Tcpdump Output

Purpose Verify that the link can be a member of both the IPv4 unicast topology and the IPv6 unicast topology.

Action user@R1> monitor traffic detail interface so-1/2/0.0
[...]

```
15:52:35.719540 In IS-IS, length 82
p2p IIH, hlen: 20, v: 1, pdu-v: 1, sys-id-len: 6 (0), max-area: 3 (0)
source-id: 0192.0168.0002, holding time: 27s, Flags: [Level 1, Level
2]
circuit-id: 0x01, PDU length: 82
Point-to-point Adjacency State TLV #240, length: 15
Adjacency State: Up (0)
Extended Local circuit-ID: 0x00000054
Neighbor System-ID: 0192.0168.0001
Neighbor Extended Local circuit-ID: 0x00000043
Protocols supported TLV #129, length: 2
NLPID(s): IPv4 (0xcc), IPv6 (0x8e)
IPv4 Interface address(es) TLV #132, length: 4
IPv4 interface address: 10.0.0.2
IPv6 Interface address(es) TLV #232, length: 16
IPv6 interface address: fe80::2a0:a514:0:24c
Area address(es) TLV #1, length: 4
Area address (length: 3): 49.0002
Restart Signaling TLV #211, length: 3
Flags [none], Remaining holding time 0s
Multi Topology TLV #229, length: 4
  IPv4 unicast Topology (0x000), Flags: [none]
  IPv6 unicast Topology (0x002), Flags: [none]
```

Meaning The IS-IS hello (IIH) packet shows that IPv4 and IPv6 are supported. The hello packet lists valid IPv4 and IPv6 addresses, and therefore the routing device can create valid next-hop entries. The supported protocols are listed in the multitopology TLV #229.

Related Documentation • [Example: Configuring IS-IS Dual Stacking of IPv4 and IPv6 Unicast Addresses on page 124](#)

CHAPTER 8

IS-IS Link and Node Protection

- [Example: Configuring Link and Node Protection for IS-IS Routes on page 141](#)

Example: Configuring Link and Node Protection for IS-IS Routes

- [Understanding Loop-Free Alternate Routes for IS-IS on page 141](#)
- [Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 145](#)

Understanding Loop-Free Alternate Routes for IS-IS

In Junos OS Release 9.5 and later, support for IS-IS loop-free alternate routes enables IP fast-reroute capability for IS-IS. Junos OS precomputes loop-free backup routes for all IS-IS routes. These backup routes are preinstalled in the Packet Forwarding Engine, which performs a local repair and implements the backup path when the link for a primary next hop for a particular route is no longer available. With local repair, the Packet Forwarding Engine can correct a path failure before it receives recomputed paths from the Routing Engine. Local repair reduces the amount of time needed to reroute traffic to less than 50 milliseconds. In contrast, global repair can take up to 800 milliseconds to compute a new route. Local repair and global repair are thus complementary. Local repair enables traffic to continue to be routed using a backup path until global repair is able to calculate a new route.

A loop-free path is one that does not forward traffic back through the routing device to reach a given destination. That is, a neighbor whose shortest path to the destination traverses the routing device is not used as a backup route to that destination. To determine loop-free alternate paths for IS-IS routes, Junos OS runs shortest-path-first (SPF) calculations on each one-hop neighbor. You can enable support for alternate loop-free routes on any IS-IS interface. Because it is common practice to enable LDP on an interface for which IS-IS is already enabled, this feature also provides support for LDP label-switched paths (LSPs).



NOTE: If you enable support for alternate loop-free routes on an interface configured for both LDP and IS-IS, you can use the `traceroute` command to trace the active path to the primary next hop.


The level of backup coverage available through IS-IS routes depends on the actual network topology and is typically less than 100 percent for all destinations on any given routing device. You can extend backup coverage to include RSVP LSPs.

Junos OS provides two mechanisms for route redundancy for IS-IS through alternate loop-free routes: link protection and node-link protection. When you enable link protection or node-link protection on an IS-IS interface, Junos OS creates a single alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection offers per-link traffic protection. Use link protection when you assume that only a single link might become unavailable but that the neighboring node on the primary path would still be available through another interface.

Node-link protection establishes an alternate path through a different routing device altogether. Use node-link protection when you assume that access to a node is lost when a link is no longer available. As a result, Junos OS calculates a backup path that avoids the primary next-hop routing device. In Junos OS Release 9.4 and earlier, only the RSVP protocol supports Packet Forwarding Engine local repair and fast reroute as well as link protection and node protection.

In [Figure 20 on page 143](#), Case 2 shows how link protection allows source Router A to switch to Link B when the primary next hop Link A to destination Router C fails. However, if Router B fails, Link B also fails, and the protected Link A is lost. If node-link protection is enabled, Router A is able to switch to Link D on Router D and bypass the failed Router B altogether. As shown in Case 1, with node-link protection enabled, Router A has a node-link protection alternate path available through Router D to destination Router C. That means that if Router B fails, Router A can still reach Router C because the path from Router A to Link D remains available as an alternate backup path.

Figure 20: Link Protection and Node-Link Protection Comparison for IS-IS Routes



The Junos OS implementation of support for loop-free alternate paths for IS-IS routes is based on the following standards:

- RFC 5286, *Basic Specification for IP Fast-Reroute: Loop-free Alternates*
- RFC 5714, *IP Fast Reroute Framework*

Configuring Link Protection for IS-IS

You can configure link protection on any interface for which IS-IS is enabled. When you enable link protection, Junos OS creates one alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection assumes that only a single link becomes unavailable but that the neighboring node would still be available through another interface.



NOTE: You must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

To enable link protection, include the **link-protection** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```
[edit]
protocols {
  isis {
    interface interface-name {
```

```

    link-protection;
  }
}

```

Configuring Node-Link Protection for IS-IS

You can configure node-link protection on any interface for which IS-IS is enabled. Node-link protection establishes an alternate path through a different routing device altogether for all destination routes that traverse a protected interface. Node-link protection assumes that the entire routing device, or node, has failed. Junos OS therefore calculates a backup path that avoids the primary next-hop routing device.



NOTE: You must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

To enable node-link protection, include the **node-link-protection** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```

[edit]
protocols {
  isis {
    interface interface-name {
      node-link-protection;
    }
  }
}

```

Excluding an IS-IS Interface as a Backup for Protected Interfaces

By default, all IS-IS interfaces that belong to the master instance or a specific routing instance are eligible as backup interfaces for protected interfaces. You can specify that any IS-IS interface be excluded from functioning as a backup interface to protected interfaces. To exclude an IS-IS interface as a backup interface, include the **no-eligible-backup** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```

[edit]
protocols {
  isis {
    interface interface-name {
      no-eligible-backup;
    }
  }
}

```

Configuring RSVP Label-Switched Paths as Backup Paths for IS-IS

Relying on the shortest-path-first (SPF) calculation of backup paths for one-hop neighbors might result in less than 100 percent backup coverage for a specific network topology. You can enhance coverage of IS-IS and LDP label-switched paths (LSPs) by configuring RSVP LSPs as backup paths. To configure a specific RSVP LSP as a backup

path, include the **backup** statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level:

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      backup;
      to ip-address;
    }
  }
}
```

When configuring an LSP, you must specify the IP address of the egress routing device with the **to** statement. For detailed information about configuring LSPs and RSVP, see the *RSVP Feature Guide for Routing Devices*.

Using Operational Mode Commands to Monitor Protected IS-IS Routes

You can issue operational mode commands that provide more details about your link-protected and node-link-protected IS-IS routes. The following guidelines explain the type of information available from the output of each command:

- **show isis backup label-switched-path**—Displays which MPLS LSPs have been designated as backup paths and the current status of those LSPs.
- **show isis backup spf results**—Displays SPF calculations for each neighbor for a given destination. Indicates whether a specific interface or node has been designated as a backup path and why. Use the **no-coverage** option to display only those nodes that do not have backup coverage.
- **show isis backup coverage**—Displays the percentage of nodes and prefixes for each type of address family that is protected.
- **show isis interface detail**—Displays the type of protection (link or node-link) applied to each protected interface.

Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN

Node-link protection establishes an alternate path through a different routing device. Use node-link protection when you assume that access to a node is lost when a link is no longer available. Junos OS calculates a backup path that avoids the primary next-hop routing device.

- [Requirements on page 145](#)
- [Overview on page 146](#)
- [Configuration on page 146](#)
- [Verification on page 153](#)

Requirements

This example requires Junos OS Release 9.5 or later.

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


Overview

In this example, core-facing interfaces are enabled for IS-IS Level 2, LDP, and RSVP. Node-link protection is enabled on all the core-facing interfaces, which means that if the primary next hop for any destination that traverses the interfaces becomes unavailable, Junos OS uses a backup link that avoids the next-hop router altogether if necessary.

You also need to configure a routing policy that requires all traffic to use per-packet load balancing in order to enable Packet Forwarding Engine local repair. With local repair, the Packet Forwarding Engine can correct a path failure and implement a backup loop-free alternate route before it receives recomputed paths from the Routing Engine.

Figure 21 on page 146 shows the topology used in this example.

Figure 21: IS-IS Node-Link Protection Topology



On Device PE1, an RSVP LSP is configured as a backup path for IS-IS. Relying on the shortest-path-first (SPF) calculation of backup paths for one-hop neighbors might result in less than 100 percent backup coverage for a specific network topology. You can enhance coverage of IS-IS and LDP LSPs by configuring RSVP LSPs as backup paths. To configure a specific RSVP LSP as a backup path, include the **backup** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level.

“CLI Quick Configuration” on page 146 shows the configuration for all of the devices in Figure 21 on page 146. The section “Step-by-Step Procedure” on page 150 describes the steps on Device P1.

Configuration

CLI Quick Configuration	To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.
Device CE1	<pre>set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30 set interfaces lo0 unit 0 family inet address 10.255.1.1/32</pre>
Device PE1	<pre>set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30 set interfaces fe-1/2/0 unit 0 family iso</pre>


```

set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.5/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.21/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.2.2/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0202.00
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface fe-1/2/1.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path to-p2 backup
set protocols mpls label-switched-path to-p2 to 10.255.4.4
set protocols mpls label-switched-path to-p2 ldp-tunneling
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface fe-1/2/1.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols bgp group l3vpn type internal
set protocols bgp group l3vpn local-address 10.255.2.2
set protocols bgp group l3vpn family inet-vpn unicast
set protocols bgp group l3vpn peer-as 65534
set protocols bgp group l3vpn local-as 65534
set protocols bgp group l3vpn neighbor 10.255.5.5
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface fe-1/2/1.0
set protocols ldp interface fe-1/2/2.0
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface fe-1/2/0.0
set routing-instances VPN-A route-distinguisher 65534:1234
set routing-instances VPN-A vrf-target target:65534:1234
set routing-instances VPN-A routing-options static route 10.255.1.1/32 next-hop 10.0.0.1
set routing-options autonomous-system 65534
set routing-options forwarding-table export ecmp

```

```

Device P1
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.6/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.9/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.25/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.3.3/32

```

```
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0303.00
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-options forwarding-table export ecmp
```

Device P2

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.10/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.13/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.4.4/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0404.00
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-options forwarding-table export ecmp
```

Device P3

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.22/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.26/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.30/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.7.7/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0707.00
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
```

```

set protocols mpls interface fxp0.0 disable
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-options forwarding-table export ecmp

```

```

Device PE2
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.14/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.29/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.5.5/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0505.00
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols bgp group l3vpn type internal
set protocols bgp group l3vpn local-address 10.255.5.5
set protocols bgp group l3vpn family inet-vpn unicast
set protocols bgp group l3vpn peer-as 65534
set protocols bgp group l3vpn local-as 65534
set protocols bgp group l3vpn neighbor 10.255.2.2
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface fe-1/2/0.0
set protocols ldp interface fe-1/2/2.0
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface fe-1/2/1.0
set routing-instances VPN-A route-distinguisher 65534:1234
set routing-instances VPN-A vrf-target target:65534:1234
set routing-instances VPN-A routing-options static route 10.255.1.1/32 next-hop 10.0.0.18
set routing-instances VPN-A routing-options autonomous-system 65534
set routing-instances VPN-A routing-options forwarding-table export ecmp

```

Device CE2 `set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.18/30`
`set interfaces lo0 unit 0 family inet address 10.255.6.6/32`

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

To configure multi-level IS-IS:

1. Configure the interfaces.

Enable IS-IS and MPLS.

```
[edit interfaces]
user@P1# set fe-1/2/0 unit 0 family inet address 10.0.0.6/30
user@P1# set fe-1/2/0 unit 0 family iso
user@P1# set fe-1/2/0 unit 0 family mpls
user@P1# set fe-1/2/1 unit 0 family inet address 10.0.0.9/30
user@P1# set fe-1/2/1 unit 0 family iso
user@P1# set fe-1/2/1 unit 0 family mpls
user@P1# set fe-1/2/2 unit 0 family inet address 10.0.0.25/30
user@P1# set fe-1/2/2 unit 0 family iso
user@P1# set fe-1/2/2 unit 0 family mpls
user@P1# set lo0 unit 0 family inet address 10.255.3.3/32
user@P1# set lo0 unit 0 family iso address 49.0001.0010.0000.0303.00
```

2. Configure the IS-IS interfaces for Level 2.

```
[edit protocols]
user@P1# set isis interface all level 2 metric 10
user@P1# set isis interface all level 1 disable
user@P1# set isis interface fxp0.0 disable
user@P1# set isis interface lo0.0 level 2 metric 0
```

3. Enable IS-IS node-link protection, which also automatically extends backup coverage to all LDP LSPs.

```
[edit protocols]
user@P1# set isis interface all node-link-protection
```

4. (Optional) Configure a 1000-millisecond time interval between the detection of a topology change and when the SPF algorithm runs.

```
[edit protocols]
user@P1# set isis spf-options delay 1000
```

5. Configure MPLS to use both RSVP and LDP label-switched paths (LSPs).

```
[edit protocols]
user@P1# set mpls interface all
user@P1# set mpls interface fxp0.0 disable
user@P1# set rsvp interface all
user@P1# set rsvp interface fxp0.0 disable
user@P1# set ldp interface all
user@P1# set ldp interface fxp0.0 disable
```

6. (Optional) For LDP, enable forwarding equivalence class (FEC) deaggregation, which results in faster global convergence.

```
[edit protocols]
user@P1# set ldp deaggregate
```

7. To enable Packet Forwarding Engine local repair, establish a policy that forces the routing protocol process to install all the next hops for a given route.

This policy ensures that the backup route is installed in the forwarding table used by the Packet Forwarding Engine to forward traffic to a given destination.

```
[edit policy-options policy-statement ecmp term 1]
user@P1# set then load-balance per-packet
```

8. Apply the policy to the forwarding table of the local router with the **export** statement.

```
[edit routing-options forwarding-table]
user@P1# set export ecmp
```

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

```
user@P1# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.6/30;
    }
    family iso;
    family mpls;
  }
}
fe-1/2/1 {
  unit 0 {
    family inet {
      address 10.0.0.9/30;
    }
    family iso;
    family mpls;
  }
}
fe-1/2/2 {
  unit 0 {
    family inet {
      address 10.0.0.25/30;
    }
    family iso;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.3.3/32;
    }
    family iso {
```

```
        address 49.0001.0010.0000.0303.00;
    }
}

user@P1# show protocols
rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
isis {
    spf-options delay 1000;
    interface all {
        node-link-protection;
        level 2 metric 10;
        level 1 disable;
    }
    interface fxp0.0 {
        disable;
    }
    interface lo0.0 {
        level 2 metric 0;
    }
}
ldp {
    deaggregate;
    interface all;
    interface fxp0.0 {
        disable;
    }
}

user@P1# show policy-options
policy-statement ecmp {
    term 1 {
        then {
            load-balance per-packet;
        }
    }
}

user@P1# show routing-options
forwarding-table {
    export ecmp;
}
```

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

Verification

Confirm that the configuration is working properly.

- [Checking the MPLS LSP Backup Path on page 153](#)
- [Checking Which Next-Hop Neighbors Are Designated as Backup Paths to the Destination Node on page 153](#)
- [Checking the Backup Coverage on page 154](#)
- [Checking the Type of Protection Configured on page 155](#)

Checking the MPLS LSP Backup Path

Purpose Display information about the MPLS label-switched-paths (LSPs) designated as the backup route for the IS-IS routes.

Action On Device PE1, from operational mode, enter the **show isis backup label-switched-path** command.

```
user@PE1> show isis backup label-switched-path
Backup MPLS LSPs:
to-p2, Egress: 10.255.4.4, Status: up, Last change: 01:17:45
TE-metric: 19, Metric: 0, Refcount: 1
```

Meaning The output shows that the backup path is up and operational.

Checking Which Next-Hop Neighbors Are Designated as Backup Paths to the Destination Node

Purpose Display SPF calculations for each neighbor for a given destination.

Action On Device PE1, from operational mode, enter the **show isis backup spf results** command.

```
user@PE1> show isis backup spf results

IS-IS level 1 SPF results:
0 nodes

IS-IS level 2 SPF results:
PE2.00
Primary next-hop: fe-1/2/2.0, IPV4, P3, SNPA: 0:5:85:8f:c8:bd
Root: P2, Root Metric: 20, Metric: 10, Root Preference: 0x0
track-item: P2.00-00
Eligible, Backup next-hop: fe-1/2/1.0, LSP, to-p2
Root: P3, Root Metric: 10, Metric: 10, Root Preference: 0x0
Not eligible, Reason: Interface is already covered
Root: P1, Root Metric: 10, Metric: 20, Root Preference: 0x0
track-item: P3.00-00
Not eligible, Reason: Interface is already covered
P2.00
Primary next-hop: fe-1/2/1.0, IPV4, P1, SNPA: 0:5:85:8f:c8:bd
Root: P2, Root Metric: 20, Metric: 0, Root Preference: 0x0
track-item: P2.00-00
Not eligible, Reason: Primary next-hop link fate sharing
Root: P1, Root Metric: 10, Metric: 10, Root Preference: 0x0
Not eligible, Reason: Primary next-hop link fate sharing
Root: P3, Root Metric: 10, Metric: 20, Root Preference: 0x0
```

```

        track-item: P1.00-00
        Not eligible, Reason: Primary next-hop node fate sharing
P3.00
    Primary next-hop: fe-1/2/2.0, IPV4, P3, SNPA: 0:5:85:8f:c8:bd
    Root: P2, Root Metric: 20, Metric: 20, Root Preference: 0x0
    track-item: P3.00-00
    track-item: P2.00-00
    track-item: P1.00-00
    Eligible, Backup next-hop: fe-1/2/1.0, LSP, to-p2
    Root: P3, Root Metric: 10, Metric: 0, Root Preference: 0x0
    Not eligible, Reason: Interface is already covered
    Root: P1, Root Metric: 10, Metric: 10, Root Preference: 0x0
    track-item: P3.00-00
    Not eligible, Reason: Interface is already covered
P1.00
    Primary next-hop: fe-1/2/1.0, IPV4, P1, SNPA: 0:5:85:8f:c8:bd
    Root: P2, Root Metric: 20, Metric: 10, Root Preference: 0x0
    track-item: P2.00-00
    track-item: P1.00-00
    Not eligible, Reason: Primary next-hop link fate sharing
    Root: P1, Root Metric: 10, Metric: 0, Root Preference: 0x0
    Not eligible, Reason: Primary next-hop link fate sharing
    Root: P3, Root Metric: 10, Metric: 10, Root Preference: 0x0
    track-item: P1.00-00
    Eligible, Backup next-hop: fe-1/2/2.0, IPV4, P3, SNPA: 0:5:85:8f:c8:bd
4 nodes

```

Meaning The output indicates whether a specific interface or node has been designated as a backup path and why.

Checking the Backup Coverage

Purpose Check the percentage of protected nodes and prefixes.

Action From operational mode, enter the **show isis backup coverage** command.

```
user@PE1> show isis backup coverage
```

```
Backup Coverage:
```

Topology	Level	Node	IPv4	IPv6	CLNS
IPV4 Unicast	1	0.00%	0.00%	0.00%	0.00%
IPV4 Unicast	2	75.00%	87.50%	0.00%	0.00%

```
user@P1> show isis backup coverage
```

```
Backup Coverage:
```

Topology	Level	Node	IPv4	IPv6	CLNS
IPV4 Unicast	1	0.00%	0.00%	0.00%	0.00%
IPV4 Unicast	2	75.00%	71.43%	0.00%	0.00%

```
user@P2> show isis backup coverage
```

```
Backup Coverage:
```

Topology	Level	Node	IPv4	IPv6	CLNS
IPV4 Unicast	1	0.00%	0.00%	0.00%	0.00%
IPV4 Unicast	2	50.00%	37.50%	0.00%	0.00%

```
user@P3> show isis backup coverage
```

```
Backup Coverage:
```

Topology	Level	Node	IPv4	IPv6	CLNS
IPV4 Unicast	1	0.00%	0.00%	0.00%	0.00%
IPV4 Unicast	2	75.00%	71.43%	0.00%	0.00%


```

user@PE2> show isis backup coverage
Backup Coverage:
Topology      Level  Node   IPv4   IPv6   CLNS
IPv4 Unicast  1     0.00%  0.00%  0.00%  0.00%
IPv4 Unicast  2     50.00% 37.50%  0.00%  0.00%

```

Meaning The level of backup coverage available through IS-IS routes depends on the actual network topology and is typically less than 100 percent for all destinations on any given routing device. You can extend backup coverage to include RSVP LSPs.

Checking the Type of Protection Configured

Purpose On all nodes in the IS-IS domain, check the percentage of protected nodes and prefixes.

Action From operational mode, enter the **show isis interface detail** command.

```

user@PE1> show isis interface detail

IS-IS interface database:
lo0.0
  Index: 76, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Adjacency advertisement: Advertise
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1           0       64      0 Passive
    2           0       64      0 Passive
fe-1/2/2.0
  Index: 79, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 10 s
  Adjacency advertisement: Advertise
  Protection Type: Node Link
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    2           1       64      10    9.000    27 P3.03 (not us)
fe-1/2/1.0
  Index: 77, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 10 s
  Adjacency advertisement: Advertise
  Protection Type: Node Link
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    2           1       64      10    9.000    27 P1.02 (not us)

```

Meaning The output shows that node-link protection is configured on the interfaces.

Related Documentation

- [Example: Configuring BFD for IS-IS on page 91](#)

CHAPTER 9

IS-IS Traffic Engineering

- [IS-IS Extensions to Support Traffic Engineering on page 157](#)
- [Example: Enabling IS-IS Traffic Engineering Support on page 158](#)
- [Example: Advertising Label-Switched Paths into IS-IS on page 172](#)
- [Example: Enabling Wide IS-IS Metrics for Traffic Engineering on page 180](#)
- [Example: Configuring Synchronization Between IS-IS and LDP on page 183](#)

IS-IS Extensions to Support Traffic Engineering

To help provide traffic engineering and MPLS with information about network topology and loading, extensions have been added to the Junos OS implementation of IS-IS. Specifically, IS-IS supports new type, length, and value (TLV) tuples that specify link attributes. These TLVs are included in the IS-IS link-state PDUs. The link-attribute information is used to populate the traffic engineering database, which is used by the Constrained Shortest Path First (CSPF) algorithm to compute the paths that MPLS label-switched paths (LSPs) take. This path information is used by RSVP to set up LSPs and reserve bandwidth for them.



NOTE: Whenever possible, use IS-IS interior gateway protocol (IGP) shortcuts instead of traffic engineering shortcuts.

The traffic engineering extensions are defined in RFC 5305, *IS-IS Extensions for Traffic Engineering*.

IS-IS IGP Shortcuts

In IS-IS, you can configure shortcuts, which allow IS-IS to use an LSP as the next hop as if it were a subinterface from the ingress routing device to the egress routing device. The address specified in the **to** statement at the **[edit protocols mpls label-switched-path lsp-path-name]** hierarchy level must match the router ID of the egress routing device for the LSP to function as a direct link to the egress routing device and to be used as input to IS-IS SPF calculations. When used in this way, LSPs are no different than Asynchronous Transfer Mode (ATM) and Frame Relay virtual circuits (VCs), except that LSPs carry only IPv4 traffic.

Example: Enabling IS-IS Traffic Engineering Support


- [IGP Shortcuts on page 158](#)
- [Example: Enabling IS-IS Traffic Engineering Support on page 159](#)

IGP Shortcuts

Link-state protocols, such as OSPF and IS-IS, use the shortest-path-first (SPF) algorithm to compute the shortest-path tree to all nodes in the network. The results of such computations can be represented by the destination node, next-hop address, and output interface, where the output interface is a physical interface. Label-switched paths (LSPs) can be used to augment the SPF algorithm, for the purposes of resolving BGP next hops. On the node performing the calculations, LSPs appear to be logical interfaces directly connected to remote nodes in the network. If you configure the interior gateway protocol (IGP) to treat LSPs the same as a physical interface and use the LSPs as a potential output interface, the SPF computation results are represented by the destination node and output LSP, effectively using the LSP as a shortcut through the network to the destination.


As an illustration, begin with a typical SPF tree (see [Figure 22 on page 158](#)).

Figure 22: Typical SPF Tree, Sourced from Router A



If an LSP connects Router A to Router D and if IGP shortcuts are enabled on Router A, you might have the SPF tree shown in [Figure 23 on page 158](#).

Figure 23: Modified SPF Tree, Using LSP A–D as a Shortcut



Note that Router D is now reachable through LSP A–D.

When computing the shortest path to reach Router D, Router A has two choices:

- Use IGP path A–B–D.


- Use LSP A–D.

Router A decides between the two choices by comparing the IGP metrics for path A–B–D with the LSP metrics for LSP A–D. If the IGP metric is lower, path A–B–D is chosen (Figure 22 on page 158). If the LSP metric is lower, LSP A–D is used (Figure 23 on page 158). If both metrics are equal, LSP A–D is chosen because LSPs are preferred over IGP paths.

Note that Routers E and F are also reachable through LSP A–D, because they are downstream from Router D in the SPF tree.

Assuming that another LSP connects Router A to Router E, you might have the SPF tree shown in Figure 24 on page 159.

Figure 24: Modified SPF Tree, Using LSP A–D and LSP A–E as Shortcuts



Example: Enabling IS-IS Traffic Engineering Support

This example shows how to configure IS-IS so that it uses label-switched paths as shortcuts.

- [Requirements on page 159](#)
- [Overview on page 159](#)
- [Configuration on page 161](#)
- [Verification on page 167](#)

Requirements

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


Overview

MPLS traffic engineering maps certain data flows to established label-switched paths (LSPs) rather than to data links calculated by the interior gateway protocol (IGP) to be part of the best (shortest) path. Fundamental to this function is the determination of what traffic is to be mapped to an LSP. Traffic is mapped to an LSP at the tunnel's ingress label switching router (LSR) by designating the egress LSR as the next-hop router for certain destination prefixes.

It is important to understand that the LSP does not constitute an entire route to a destination. Rather, the LSP is a next-hop segment of the route. Therefore, packets can only be mapped to an LSP if the egress LSR is considered to be a feasible next-hop candidate during the route resolution process.

Figure 25 on page 160 shows the topology used in this example.

Figure 25: IS-IS Shortcuts Topology



In this example, Device C has an external BGP (EBGP) peer session with Device G in autonomous system (AS) 65520. In order to enable its internal BGP (IBGP) peers to access subnets in AS 65520, Device G runs IS-IS passively on its interface connecting to Device C. IS-IS has information about the external subnets and enters routes to these subnets in the inet.0 routing table. BGP, when resolving the next-hop addresses of AS-external routes, uses the IGP route.



TIP: An alternative to passively running IS-IS on the interface would be to use a next-hop self policy.

Device A has an LSP to Device C. The path is configured to always go through Device E, rather than going through Device B.

Interior gateway protocol (IGP) shortcuts, also called traffic-engineering shortcuts, provide a tool by which the link-state IGP (OSPF or IS-IS) in an AS can consider an LSP in its shortest-path-first (SPF) calculations. If using passive external interfaces, the IGP views an LSP as a single data link toward the destinations beyond the LSP egress device.

When you use **traffic-engineering bgp** (which is the default) and IGP shortcuts, the traffic engineering solution is used for BGP AS-external route resolution only. However, traffic to AS-internal destinations can also be mapped to LSPs. To accomplish this, **traffic-engineering bgp-igp** is enabled. Thus, RSVP installs the MPLS prefixes into the inet.0 table rather than the inet.3 table. As a result, the MPLS LSPs are installed in the forwarding table.

This approach finds practical application whenever heavy traffic is routed to specific destinations within an AS, such as server farms.

An important point about IGP shortcuts, whether used alone or in conjunction with traffic-engineering BGP-IGP, is that IGP adjacencies are never formed across the LSPs. The IGP sees the LSP as a single data link, but does not view the egress router as a potential peer and does not forward hello messages across the LSP. Also, RSVP messages are never forwarded over LSPs, preventing the possibility of an LSP being inadvertently built within another LSP.

[“CLI Quick Configuration” on page 161](#) shows the configuration for all of the devices in [Figure 25 on page 160](#). The section [“Step-by-Step Procedure” on page 164](#) describes the steps on Device A.

Configuration

CLI Quick Configuration	To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.
Device A	<pre> set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30 set interfaces fe-1/2/0 unit 0 family iso set interfaces fe-1/2/0 unit 0 family mpls set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.5/30 set interfaces fe-1/2/1 unit 0 family iso set interfaces fe-1/2/1 unit 0 family mpls set interfaces lo0 unit 0 family inet address 192.168.0.1/32 set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00 set protocols rsvp interface lo0.0 set protocols rsvp interface fe-1/2/0.0 set protocols rsvp interface fe-1/2/1.0 set protocols mpls traffic-engineering bgp-igp set protocols mpls label-switched-path test_path to 192.168.0.3 set protocols mpls label-switched-path test_path no-cspf set protocols mpls label-switched-path test_path primary through_E set protocols mpls path through_E 192.168.0.5 strict set protocols mpls interface fe-1/2/0.0 set protocols mpls interface fe-1/2/1.0 set protocols bgp group int type internal set protocols bgp group int local-address 192.168.0.1 set protocols bgp group int neighbor 192.168.0.5 set protocols bgp group int neighbor 192.168.0.6 set protocols bgp group int neighbor 192.168.0.2 set protocols bgp group int neighbor 192.168.0.3 set protocols isis traffic-engineering family inet shortcuts set protocols isis interface fe-1/2/0.0 level 1 disable set protocols isis interface fe-1/2/1.0 level 1 disable set protocols isis interface lo0.0 set routing-options router-id 192.168.0.1 set routing-options autonomous-system 1 </pre>
Device B	<pre> set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.6/30 set interfaces fe-1/2/0 unit 0 family iso set interfaces fe-1/2/0 unit 0 family mpls </pre>

```

set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.9/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/1.0
set protocols rsvp interface lo0.0
set protocols mpls interface fe-1/2/1.0
set protocols mpls interface fe-1/2/0.0
set protocols bgp group int type internal
set protocols bgp group int local-address 192.168.0.2
set protocols bgp group int neighbor 192.168.0.6
set protocols bgp group int neighbor 192.168.0.5
set protocols bgp group int neighbor 192.168.0.1
set protocols bgp group int neighbor 192.168.0.3
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0
set routing-options router-id 192.168.0.2
set routing-options autonomous-system 1

```

Device C

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.10/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.13/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.25/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/3 unit 0 family inet address 10.0.0.29/30
set interfaces fe-1/2/3 unit 0 family iso
set interfaces fe-1/2/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0003.00
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fe-1/2/3.0
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/3.0
set protocols bgp group int type internal
set protocols bgp group int local-address 192.168.0.3
set protocols bgp group int neighbor 192.168.0.6
set protocols bgp group int neighbor 192.168.0.5
set protocols bgp group int neighbor 192.168.0.1
set protocols bgp group int neighbor 192.168.0.2
set protocols bgp group external-peers type external
set protocols bgp group external-peers export send-some-isis
set protocols bgp group external-peers peer-as 2
set protocols bgp group external-peers neighbor 10.0.0.26
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface fe-1/2/2.0 level 1 disable
set protocols isis interface fe-1/2/2.0 level 2 passive
set protocols isis interface fe-1/2/3.0 level 1 disable
set protocols isis interface lo0.0
set policy-options policy-statement send-some-isis term 1 from protocol isis

```



```

set policy-options policy-statement send-some-isis term 1 from route-filter 10.0.0.0/24
  orlonger
set policy-options policy-statement send-some-isis term 1 from route-filter 192.168.0.0/24
  orlonger
set policy-options policy-statement send-some-isis term 1 then accept
set routing-options router-id 192.168.0.3
set routing-options autonomous-system 1

```

Device D

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.14/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0004.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0
set routing-options router-id 192.168.0.4
set routing-options autonomous-system 1

```

Device E

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.0.5/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0005.00
set protocols rsvp interface lo0.0
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/1.0
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/1.0
set protocols bgp group int type internal
set protocols bgp group int local-address 192.168.0.5
set protocols bgp group int neighbor 192.168.0.1
set protocols bgp group int neighbor 192.168.0.6
set protocols bgp group int neighbor 192.168.0.2
set protocols bgp group int neighbor 192.168.0.3
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0
set routing-options router-id 192.168.0.5
set routing-options autonomous-system 1

```

Device F

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.18/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.30/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.0.6/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0006.00
set protocols rsvp interface lo0.0
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/1.0
set protocols rsvp interface fe-1/2/2.0

```

```

set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/1.0
set protocols mpls interface fe-1/2/2.0
set protocols bgp group int type internal
set protocols bgp group int local-address 192.168.0.6
set protocols bgp group int neighbor 192.168.0.1
set protocols bgp group int neighbor 192.168.0.5
set protocols bgp group int neighbor 192.168.0.2
set protocols bgp group int neighbor 192.168.0.3
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface fe-1/2/2.0 level 1 disable
set protocols isis interface lo0.0
set routing-options router-id 192.168.0.6
set routing-options autonomous-system 1

```

Device G

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.26/30
set interfaces lo0 unit 0 family inet address 192.168.0.7/32 primary
set interfaces lo0 unit 0 family inet address 10.2.1.1/32
set interfaces lo0 unit 0 family inet address 10.3.1.1/32
set protocols bgp group external-peers type external
set protocols bgp group external-peers export statics
set protocols bgp group external-peers export send-directs
set protocols bgp group external-peers peer-as 1
set protocols bgp group external-peers neighbor 10.0.0.25
set policy-options policy-statement statics from protocol static
set policy-options policy-statement statics then accept
set policy-options policy-statement send-directs term 1 from protocol direct
set policy-options policy-statement send-directs term 1 then accept
set routing-options static route 10.2.0.0/32 reject
set routing-options static route 10.2.0.0/32 install
set routing-options static route 10.3.0.0/32 reject
set routing-options static route 10.3.0.0/32 install
set routing-options router-id 192.168.0.7
set routing-options autonomous-system 2

```

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

To configure IS-IS traffic-engineering shortcuts:

1. Configure the interfaces.

```

[edit interfaces]
user@A# set fe-1/2/0 unit 0 family inet address 10.0.0.1/30
user@A# set fe-1/2/0 unit 0 family iso
user@A# set fe-1/2/0 unit 0 family mpls
user@A# set fe-1/2/1 unit 0 family inet address 10.0.0.5/30
user@A# set fe-1/2/1 unit 0 family iso
user@A# set fe-1/2/1 unit 0 family mpls
user@A# set lo0 unit 0 family inet address 192.168.0.1/32
user@A# set lo0 unit 0 family iso address 49.0002.0192.0168.0001.00

```

2. Enable a signaling protocol on the interfaces.

```
[edit protocols rsvp]
user@A# set interface lo0.0
user@A# set interface fe-1/2/0.0
user@A# set interface fe-1/2/1.0
```

3. Enable MPLS on the interfaces.

```
[edit protocols mpls]
user@A# set interface fe-1/2/0.0
user@A# set interface fe-1/2/1.0
```

4. Configure the label-switched path.

A single LSP, named `test_path`, is configured from Device A to Device C. The LSP explicit route object (ERO) is specified to use a strict hop through Device E, so that the LSP takes a different path from the OSPF shortest path of A–B–C. The LSP is signaled using RSVP, but no CSPF is running.

```
[edit protocols mpls]
user@A# set label-switched-path test_path to 192.168.0.3
user@A# set label-switched-path test_path no-cspf
user@A# set label-switched-path test_path primary through_E
user@A# set path through_E 192.168.0.5 strict
```

5. Configure traffic engineering for both BGP and IGP destinations.

When IGP shortcuts are also enabled, the IGP can use the LSP in its calculations. The results of the calculations are entered into the `inet.0` table.

```
[edit protocols mpls]
user@A# set traffic-engineering bgp-igp
```

6. Configure internal BGP (IBGP) peering among the devices.

```
[edit protocols bgp group int]
user@A# set type internal
user@A# set local-address 192.168.0.1
user@A# set neighbor 192.168.0.5
user@A# set neighbor 192.168.0.6
user@A# set neighbor 192.168.0.2
user@A# set neighbor 192.168.0.3
```

7. Enable IS-IS on the interfaces, and set the link metric.

```
[edit protocols isis]
user@A# set interface fe-1/2/0.0 level 1 disable
user@A# set interface fe-1/2/1.0 level 1 disable
user@A# set interface lo0.0
```

8. Configure IS-IS to use MPLS LSPs as next hops for the IPv4 address family.

It is only necessary to enable IGP shortcuts on the ingress router because that is the router performing the shortest-path-first (SPF) calculations.

It is important to understand how IGP shortcuts affect the protocol and routing table relationship. The IGP performs SPF calculations to subnets downstream of LSP egress points, but the results of these calculations are entered into the `inet.3` table only. At the same time, the IGP performs its traditional SPF calculations and enters the results of these calculations into the `inet.0` table. The result is that

although the IGP is making entries into the inet.3 table, BGP is still the only protocol with visibility into that table for the purposes of route resolution. Therefore, forwarding to AS-internal destinations still uses the inet.0 IGP routes, and the LSPs are only used for BGP next-hop resolution. If you want the LSPs to be used for IGP next-hop resolution, you must configure **traffic-engineering bgp-igp**.

```
[edit protocols isis]
user@A# set traffic-engineering family inet shortcuts
```

9. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@A# set router-id 192.168.0.1
user@A# set autonomous-system 1
```

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

```
user@A# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.1/30;
    }
    family iso;
    family mpls;
  }
}
fe-1/2/1 {
  unit 0
    family inet {
      address 10.0.0.5/30;
    }
    family iso;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.1/32;
    }
    family iso {
      address 49.0002.0192.0168.0001.00;
    }
  }
}

user@A# show protocols
rsvp {
  interface lo0.0;
  interface fe-1/2/0.0;
  interface fe-1/2/1.0;
}
mpls {
```

```

traffic-engineering bgp-igp;
label-switched-path test_path {
  to 192.168.0.3;
  no-cspf;
  primary through_E;
}
path through_E {
  192.168.0.5 strict;
}
interface fe-1/2/0.0;
interface fe-1/2/1.0;
}
bgp {
  group int {
    type internal;
    local-address 192.168.0.1;
    neighbor 192.168.0.5;
    neighbor 192.168.0.6;
    neighbor 192.168.0.2;
    neighbor 192.168.0.3;
  }
}
isis {
  traffic-engineering {
    family inet {
      shortcuts;
    }
  }
  interface fe-1/2/0.0 {
    level 1 disable;
  }
  interface fe-1/2/1.0 {
    level 1 disable;
  }
  interface lo0.0;
}

user@A# show routing-options
router-id 192.168.0.1;
autonomous-system 1;

```

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

Verification

Confirm that the configuration is working properly.

- [Verifying the Next Hops on page 167](#)
- [Checking the RSVP Sessions on page 169](#)
- [Checking the Paths with Different Traffic Engineering Settings on page 170](#)

Verifying the Next Hops

Purpose Verify that the MPLS LSP is used as the next hop in the expected routes.

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

```
user@A> show route
```

```
inet.0: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30      *[Direct/0] 4d 09:07:26
                 > via fe-1/2/0.0
10.0.0.1/32      *[Local/0] 4d 09:07:26
                 Local via fe-1/2/0.0
10.0.0.4/30      *[Direct/0] 4d 09:07:28
                 > via fe-1/2/1.0
10.0.0.5/32      *[Local/0] 4d 09:07:28
                 Local via fe-1/2/1.0
10.0.0.8/30      *[IS-IS/18] 01:42:24, metric 20
                 > to 10.0.0.6 via fe-1/2/1.0
10.0.0.12/30     *[IS-IS/18] 01:42:24, metric 30
                 > to 10.0.0.6 via fe-1/2/1.0
10.0.0.16/30     *[IS-IS/18] 01:42:24, metric 20
                 > to 10.0.0.2 via fe-1/2/0.0
10.0.0.20/30     *[IS-IS/18] 01:42:24, metric 30
                 > to 10.0.0.2 via fe-1/2/0.0
10.0.0.24/30     *[IS-IS/18] 01:42:24, metric 30
                 > to 10.0.0.6 via fe-1/2/1.0
10.0.0.28/30     *[IS-IS/18] 01:42:24, metric 30
                 to 10.0.0.6 via fe-1/2/1.0
                 > to 10.0.0.2 via fe-1/2/0.0
10.2.0.0/32      *[BGP/170] 02:22:30, localpref 100, from 192.168.0.3
                 AS path: 2 I, validation-state: unverified
                 > to 10.0.0.2 via fe-1/2/0.0, label-switched-path test_path
10.2.1.1/32      *[BGP/170] 02:20:23, localpref 100, from 192.168.0.3
                 AS path: 2 I, validation-state: unverified
                 > to 10.0.0.2 via fe-1/2/0.0, label-switched-path test_path
10.3.0.0/32      *[BGP/170] 02:22:30, localpref 100, from 192.168.0.3
                 AS path: 2 I, validation-state: unverified
                 > to 10.0.0.2 via fe-1/2/0.0, label-switched-path test_path
10.3.1.1/32      *[BGP/170] 02:20:23, localpref 100, from 192.168.0.3
                 AS path: 2 I, validation-state: unverified
                 > to 10.0.0.2 via fe-1/2/0.0, label-switched-path test_path
192.168.0.1/32   *[Direct/0] 4d 09:08:47
                 > via lo0.0
192.168.0.2/32   *[IS-IS/18] 01:42:24, metric 10
                 > to 10.0.0.6 via fe-1/2/1.0
192.168.0.3/32   *[IS-IS/18] 01:42:24, metric 20
                 > to 10.0.0.6 via fe-1/2/1.0
192.168.0.4/32   *[IS-IS/18] 01:42:24, metric 30
                 > to 10.0.0.6 via fe-1/2/1.0
                 to 10.0.0.2 via fe-1/2/0.0
192.168.0.5/32   *[IS-IS/18] 01:42:24, metric 10
                 > to 10.0.0.2 via fe-1/2/0.0
192.168.0.6/32   *[IS-IS/18] 01:42:24, metric 20
                 > to 10.0.0.2 via fe-1/2/0.0
192.168.0.7/32   *[BGP/170] 02:20:23, localpref 100, from 192.168.0.3
                 AS path: 2 I, validation-state: unverified
                 > to 10.0.0.2 via fe-1/2/0.0, label-switched-path test_path

inet.3: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.12/30     *[IS-IS/18] 01:41:21, metric 30
```

```

10.0.0.24/30      > to 10.0.0.2 via fe-1/2/0.0, label-switched-path test_path
                  *[IS-IS/18] 01:41:21, metric 30
10.0.0.28/30      > to 10.0.0.2 via fe-1/2/0.0, label-switched-path test_path
                  *[IS-IS/18] 01:41:21, metric 30
192.168.0.3/32    > to 10.0.0.2 via fe-1/2/0.0, label-switched-path test_path
                  *[RSVP/7/1] 01:41:21, metric 20
                  > to 10.0.0.2 via fe-1/2/0.0, label-switched-path test_path
                  [IS-IS/18] 01:41:21, metric 20
192.168.0.4/32    > to 10.0.0.2 via fe-1/2/0.0, label-switched-path test_path
                  *[IS-IS/18] 01:41:21, metric 30
                  > to 10.0.0.2 via fe-1/2/0.0, label-switched-path test_path

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0002.0192.0168.0001/72
      *[Direct/0] 4d 09:08:47
      > via lo0.0

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 4d 09:10:00, metric 1
           Receive
1          *[MPLS/0] 4d 09:10:00, metric 1
           Receive
2          *[MPLS/0] 4d 09:10:00, metric 1
           Receive
13         *[MPLS/0] 4d 09:10:00, metric 1
           Receive

```

Meaning IS-IS chooses the LSP as the shortest path to destinations downstream of the LSP egress device. Additionally, because the IGP uses the LSP to reach external subnet 10.0.0.24/30, BGP also uses the LSP in its routes to 10.2.0.0 and 10.3.0.0.

If next-hop self were used at Device C, BGP would still choose the LSP over the IGP path.

Checking the RSVP Sessions

Purpose Display information about RSVP sessions

Action From operational mode, enter the **show rsvp session brief** command.

```

user@A> show rsvp session brief
Ingress RSVP: 1 sessions
To          From          State   Rt Style Labelin Labelout LSPname
192.168.0.3 192.168.0.1 Up       0 1 FF      -    299776 test_path
Total 1 displayed, Up 1, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@E> show rsvp session brief
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

```

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 1 sessions
To          From          State   Rt  Style Labelin Labelout LSPname
192.168.0.3 192.168.0.1 Up      0  1 FF  299776  299808 test_path
Total 1 displayed, Up 1, Down 0

user@F> show rsvp session brief
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 1 sessions
To          From          State   Rt  Style Labelin Labelout LSPname
192.168.0.3 192.168.0.1 Up      0  1 FF  299808      3 test_path
Total 1 displayed, Up 1, Down 0

user@C> show rsvp session brief
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 1 sessions
To          From          State   Rt  Style Labelin Labelout LSPname
192.168.0.3 192.168.0.1 Up      0  1 FF      3      - test_path
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning On all four routing devices, the ingress and egress IP addresses of the LSP are shown. The path is shown as an ingress path at Device A, and packets forwarded on the LSP are assigned a label of 299776. At Device E, the LSP is transit, and packets arriving with a label of 299776 are given an outgoing label of 299808. The labels have significance only between neighboring label-switched routers (LSRs). Device F swaps incoming label 299808 for outgoing label 3. Device C, the egress, pops label 3 and routes the received packet by standard IP longest-match route lookup.

Checking the Paths with Different Traffic Engineering Settings

Purpose Check the paths used for IGP and BGP routes when **traffic-engineering bgp-igp** is used and when **traffic-engineering bgp** (the default) is used.

Action 1. Configure **traffic-engineering bgp**.

This removes **traffic-engineering bgp-igp** from the configuration because only one MPLS traffic engineering setting can be configured in each routing instance.

```

[edit protocols mpls]
user@A# set traffic-engineering bgp
user@A# commit

```

2. Use the **show route forwarding-table** command to check the paths when **traffic-engineering bgp** (the default) is configured.


```
user@A> show route forwarding-table destination 10.2.1.1
```

```
Routing table: default.inet
```

```
Internet:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
10.2.1.1/32	user	0	10.0.0.2	indr	262145	6	
				Push	299776	1013	2

```
fe-1/2/0.0
```

```
user@A> show route forwarding-table destination 192.168.0.3
```

```
Routing table: default.inet
```

```
Internet:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
192.168.0.3/32	user	1	10.0.0.6	ucst	938	11	fe-1/2/1.0

3. Use the **traceroute** command to check the paths when **traffic-engineering bgp** (the default) is configured.

```
user@A> traceroute 10.2.1.1
```

```
traceroute to 10.2.1.1 (10.2.1.1), 30 hops max, 40 byte packets
```

```
1 10.0.0.2 (10.0.0.2) 11.086 ms 1.587 ms 1.603 ms
   MPLS Label=299776 CoS=0 TTL=1 S=1
2 10.0.0.18 (10.0.0.18) 1.455 ms 1.477 ms 1.442 ms
   MPLS Label=299808 CoS=0 TTL=1 S=1
3 10.0.0.29 (10.0.0.29) 2.240 ms 1.045 ms 1.243 ms
4 10.2.1.1 (10.2.1.1) 1.363 ms 1.389 ms 1.374 ms
```

```
user@A> traceroute 192.168.0.3
```

```
traceroute to 192.168.0.3 (192.168.0.3), 30 hops max, 40 byte packets
```

```
1 10.0.0.6 (10.0.0.6) 1.759 ms 1.872 ms 2.281 ms
2 bb03-cc1ab-1o0.spg1ab.juniper.net (192.168.0.3) 2.119 ms 2.157 ms 1.598 ms
```

4. Configure **traffic-engineering bgp-igp**.

This removes **traffic-engineering bgp** from the configuration because only one MPLS traffic engineering setting can be configured in each routing instance.

```
[edit protocols mpls]
```

```
user@A# set traffic-engineering bgp-igp
```

```
user@A# commit
```

5. Use the **show route forwarding-table** command to check the paths when **traffic-engineering bgp-igp** is configured.

```
user@A> show route forwarding-table destination 10.2.1.1
```

```
Routing table: default.inet
```

```
Internet:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
10.2.1.1/32	user	0	10.0.0.2	indr	262145	6	
				Push	299776	1013	2

```
fe-1/2/0.0
```

```
user@A> show route forwarding-table destination 192.168.0.3
```

```
Routing table: default.inet
```

```
Internet:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
192.168.0.3/32	user	1	10.0.0.2	Push	299776	1013	8

```
fe-1/2/0.0
```

6. Use the **traceroute** command to check the paths when **traffic-engineering bgp-igp** is configured.

```
user@A> traceroute 10.2.1.1
```

```

traceroute to 10.2.1.1 (10.2.1.1), 30 hops max, 40 byte packets
 1  10.0.0.2 (10.0.0.2)  2.348 ms  1.475 ms  1.434 ms
    MPLS Label=299776 CoS=0 TTL=1 S=1
 2  10.0.0.18 (10.0.0.18)  1.507 ms  2.307 ms  1.911 ms
    MPLS Label=299808 CoS=0 TTL=1 S=1
 3  10.0.0.29 (10.0.0.29)  1.743 ms  1.645 ms  1.940 ms
 4  10.2.1.1 (10.2.1.1)  2.041 ms  1.977 ms  2.233 ms

user@A> traceroute 192.168.0.3
traceroute to 192.168.0.3 (192.168.0.3), 30 hops max, 40 byte packets
 1  10.0.0.2 (10.0.0.2)  1.721 ms  2.558 ms  2.229 ms
    MPLS Label=299776 CoS=0 TTL=1 S=1
 2  10.0.0.18 (10.0.0.18)  2.505 ms  1.462 ms  1.408 ms
    MPLS Label=299808 CoS=0 TTL=1 S=1
 3  bb03-cc1ab-1o0.spg1ab.juniper.net (192.168.0.3)  1.371 ms  1.422 ms  1.351
ms

```

Meaning When **traffic-engineering bgp** is configured, the first trace is to a destination belonging to the BGP-learned 10.2.0.0/16 prefix, and follows the LSP. The second trace is to the IS-IS-learned 192.168.0.3 route (Device C's loopback interface address), and follows the IS-IS route. These results correspond to what we observe in the forwarding table. The forwarding table is built based on routes in inet.0 only. BGP can look into inet.3 and select an LSP as the best path to the next hop of a BGP prefix, and can add a route into inet.0 utilizing that LSP. An entry is then made to the forwarding table from the inet.0 route. No other protocol, by default, can consult inet.3, and the inet.3 routes are not entered into inet.0. Therefore, the forwarding entry for 192.168.0.3 is created from the only route to that destination in inet.0: the IS-IS route.

When **traffic-engineering bgp-igp** is configured, the first trace to 10.2.1.1 continues to follow the LSP. The second trace to 192.168.0.3 also follows the LSP. These results correspond to what we observe in the forwarding table, which shows that the LSP is used for IGP next-hop resolution.

Related Documentation

- [Example: Enabling OSPF Traffic Engineering Support](#)

Example: Advertising Label-Switched Paths into IS-IS

- [Understanding Forwarding Adjacencies on page 172](#)
- [Example: Advertising Label-Switched Paths into IS-IS on page 173](#)

Understanding Forwarding Adjacencies

A forwarding adjacency is a traffic engineering label-switched path (LSP) that is configured between two nodes and that is used by the interior gateway protocol (IGP) to forward traffic.

When you set up MPLS traffic-engineering tunnels between sites, by default the IGP does not consider those tunnels for traffic forwarding. Forwarding adjacencies allow you to treat a traffic engineering LSP tunnel as a link in an IGP topology. The link is used in the shortest-path-first (SPF) algorithm and is advertised to the IGP peers. A forwarding

adjacency can be created between routing devices regardless of their location in the network.

Example: Advertising Label-Switched Paths into IS-IS

This example shows how to advertise label-switched paths (LSPs) into IS-IS as point-to-point links (sometimes referred to as forwarding adjacencies) so that the LSPs can be used in SPF calculations. The advertisement contains a local address (the **from** address of the LSP), a remote address (the **to** address of the LSP), and a metric.

- [Requirements on page 173](#)
- [Overview on page 173](#)
- [Configuration on page 173](#)
- [Verification on page 178](#)


Requirements

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

Overview

Figure 26 on page 173 shows the topology used in this example.

Figure 26: IS-IS Advertising a Label-Switched Path Topology



The example shows how to configure the LSP from Device E to Device D and then advertise this path through IS-IS. The configuration is verified by performing a traceroute operation from Device A to Device D and making sure that the LSP is used for forwarding.

“CLI Quick Configuration” on page 173 shows the configuration for all of the devices in Figure 26 on page 173. The section “Step-by-Step Procedure” on page 176 describes the steps on Device E.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

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

Device A

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.5/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set protocols rsvp interface lo0.0
set protocols rsvp interface fe-1/2/0.0
set protocols mpls interface fe-1/2/0.0
set protocols bgp group int type internal
set protocols bgp group int local-address 192.168.0.1
set protocols bgp group int neighbor 192.168.0.6
set protocols bgp group int neighbor 192.168.0.5
set protocols bgp group int neighbor 192.168.0.4
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/0.0 level 2 metric 10
set protocols isis interface fe-1/2/0.5 level 2 metric 10
set protocols isis interface fe-1/2/0.5 level 1 disable
set protocols isis interface lo0.0
set routing-options router-id 192.168.0.1
set routing-options autonomous-system 1
```

Device B

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.6/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.9/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 2 metric 10
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0
set routing-options router-id 192.168.0.2
```

Device C

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.10/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.13/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0003.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 2 metric 10
set protocols isis interface lo0.0
set routing-options router-id 192.168.0.3
```

Device D

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.14/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.22/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
```

```

set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0004.00
set protocols rsvp interface lo0.0
set protocols rsvp interface fe-1/2/1.0
set protocols mpls label-switched-path E-D to 192.168.0.5
set protocols mpls interface fe-1/2/1.0
set protocols bgp group int type internal
set protocols bgp group int local-address 192.168.0.4
set protocols bgp group int neighbor 192.168.0.5
set protocols bgp group int neighbor 192.168.0.1
set protocols bgp group int neighbor 192.168.0.6
set protocols isis interface fe-1/2/0.14 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0
set protocols isis label-switched-path E-D level 2
set routing-options router-id 192.168.0.4
set routing-options autonomous-system 1

```

Device E

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.0.5/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0005.00
set protocols rsvp interface lo0.0
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/1.0
set protocols mpls label-switched-path E-D to 192.168.0.4
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/1.0
set protocols bgp group int type internal
set protocols bgp group int local-address 192.168.0.5
set protocols bgp group int neighbor 192.168.0.6
set protocols bgp group int neighbor 192.168.0.1
set protocols bgp group int neighbor 192.168.0.4
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 2 metric 20
set protocols isis interface lo0.0
set protocols isis label-switched-path E-D level 2 metric 15
set routing-options router-id 192.168.0.5
set routing-options autonomous-system 1

```

Device F

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.18/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.21/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.0.6/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0006.00
set protocols rsvp interface lo0.0
set protocols rsvp interface fe-1/2/0.0

```

```

set protocols rsvp interface fe-1/2/1.0
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/1.0
set protocols bgp group int type internal
set protocols bgp group int local-address 192.168.0.6
set protocols bgp group int neighbor 192.168.0.1
set protocols bgp group int neighbor 192.168.0.5
set protocols bgp group int neighbor 192.168.0.4
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 2 metric 10
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0
set routing-options router-id 192.168.0.6
set routing-options autonomous-system 1

```

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

To advertise LSPs into IS-IS:

1. Configure the interfaces.

```

[edit interfaces]
user@E# set fe-1/2/0 unit 0 family inet address 10.0.0.2/30
user@E# set fe-1/2/0 unit 0 family iso
user@E# set fe-1/2/0 unit 0 family mpls
user@E# set fe-1/2/1 unit 0 family inet address 10.0.0.17/30
user@E# set fe-1/2/1 unit 0 family iso
user@E# set fe-1/2/1 unit 0 family mpls
user@E# set lo0 unit 0 family inet address 192.168.0.5/32
user@E# set lo0 unit 0 family iso address 49.0002.0192.0168.0005.00

```

2. Enable a signaling protocol on the interfaces.

```

[edit protocols rsvp]
user@E# set interface lo0.0
user@E# set interface fe-1/2/0.0
user@E# set interface fe-1/2/1.0

```

3. Enable MPLS on the interfaces.

```

[edit protocols mpls]
user@E# set interface fe-1/2/0.0
user@E# set interface fe-1/2/1.0

```

4. Configure the LSP.

Make sure that you configure the reverse LSP on the endpoint, in this case on Device D.

```

[edit protocols mpls]
user@E# set label-switched-path E-D to 192.168.0.4

```

5. Configure internal BGP (IBGP) peering among the devices that must run MPLS.

```

[edit protocols bgp group int]
user@E# set type internal
user@E# set local-address 192.168.0.5

```

```

user@E# set neighbor 192.168.0.6
user@E# set neighbor 192.168.0.1
user@E# set neighbor 192.168.0.4

```

6. Enable IS-IS on the interfaces, and set the link metric.

IS-IS Level 1 and Level 2 are enabled when you include the interface at **[edit protocols isis]**. By disabling Level 1, you are in effect creating a Level 2 IS-IS interface.

```

[edit protocols isis]
user@E# set interface fe-1/2/0.0 level 1 disable
user@E# set interface fe-1/2/1.0 level 1 disable
user@E# set interface fe-1/2/1.0 level 2 metric 20
user@E# set interface lo0.0

```

7. Advertise the LSP through IS-IS.

Make sure that you advertise the LSP on the endpoint, in this case on Device D.

```

[edit protocols isis]
user@E# set label-switched-path E-D level 2 metric 15

```

8. Configure the router ID and the autonomous system (AS) number.

```

[edit routing-options]
user@E# set router-id 192.168.0.5
user@E# set autonomous-system 1

```

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

```

user@E# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.2/30;
    }
    family iso;
    family mpls;
  }
}
fe-1/2/1 {
  unit 0 {
    family inet {
      address 10.0.0.17/30;
    }
    family iso;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.5/32;
    }
    family iso {
      address 49.0002.0192.0168.0005.00;
    }
  }
}

```

```
    }  
  }  
}  
  
user@E# show protocols  
rsvp {  
  interface lo0.0;  
  interface fe-1/2/0.0;  
  interface fe-1/2/1.0;  
}  
mpls {  
  label-switched-path E-D {  
    to 192.168.0.4;  
  }  
  interface fe-1/2/0.0;  
  interface fe-1/2/1.0;  
}  
bgp {  
  group int {  
    type internal;  
    local-address 192.168.0.5;  
    neighbor 192.168.0.6;  
    neighbor 192.168.0.1;  
    neighbor 192.168.0.4;  
  }  
}  
isis {  
  interface fe-1/2/0.0 {  
    level 1 disable;  
  }  
  interface fe-1/2/1.0 {  
    level 1 disable;  
    level 2 metric 20;  
  }  
  interface lo0.0;  
  label-switched-path E-D {  
    level 2 metric 15;  
  }  
}  
  
user@E# show routing-options  
router-id 192.168.0.5;  
autonomous-system 1;
```

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

Verification

Confirm that the configuration is working properly.

- [Verifying the IS-IS Neighbor on page 179](#)
- [Checking the IS-IS SPF Calculations on page 179](#)
- [Checking the Forwarding Path on page 180](#)

Verifying the IS-IS Neighbor

Purpose Verify that another neighbor is listed and is reachable over the LSP. The interface field indicates the name of the LSP.

Action From operational mode, enter the **show isis adjacency detail** command.

```
user@E> show isis adjacency detail
D
  Interface: E-D, Level: 2, State: One-way, Expires in 0 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 1d 00:34:58 ago
  Circuit type: 3, Speaks: IP
  Topologies: Unicast
  Restart capable: No, Adjacency advertisement: Advertise
  IP addresses: 192.168.0.4

F
  Interface: fe-1/2/1.0, Level: 2, State: Up, Expires in 7 secs
  Priority: 64, Up/Down transitions: 1, Last transition: 1d 01:16:22 ago
  Circuit type: 2, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
  Topologies: Unicast
  Restart capable: Yes, Adjacency advertisement: Advertise
  LAN id: F.02, IP addresses: 10.0.0.18

A
  Interface: fe-1/2/0.0, Level: 2, State: Up, Expires in 20 secs
  Priority: 64, Up/Down transitions: 1, Last transition: 1d 01:17:20 ago
  Circuit type: 2, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
  Topologies: Unicast
  Restart capable: Yes, Adjacency advertisement: Advertise
  LAN id: E.02, IP addresses: 10.0.0.1
```

Meaning As expected, **Interface: E-D** is shown in the output, and the state is shown as **One-way**.

Checking the IS-IS SPF Calculations

Purpose Verify that the LSP is being used in the SPF calculations.

Action From operational mode, enter the **show isis spf brief** command.

```
user@E> show isis spf brief

IS-IS level 1 SPF results:
Node      Metric  Interface  NH  Via  SNPA
E.00      0
1 nodes

IS-IS level 2 SPF results:
Node      Metric  Interface  NH  Via  SNPA
C.02      30      fe-1/2/0.0  IPV4 A  0:5:85:8f:c8:bc
C.00      25      fe-1/2/1.0  LSP E-D
D.03      25      fe-1/2/1.0  LSP E-D
D.02      25      fe-1/2/1.0  LSP E-D
F.00      20      fe-1/2/1.0  IPV4 F  0:5:85:8f:c8:bd
B.00      20      fe-1/2/0.0  IPV4 A  0:5:85:8f:c8:bc
B.02      20      fe-1/2/0.0  IPV4 A  0:5:85:8f:c8:bc
D.00      15      fe-1/2/1.0  LSP E-D
A.00      10      fe-1/2/0.0  IPV4 A  0:5:85:8f:c8:bc
```

```
E.02      10
E.00      0
  11 nodes
```

Meaning As expected, the SPF results include the LSP, E-D.

Checking the Forwarding Path

Purpose Verify that a traceroute operation from Device A to Device D uses the LSP.

Action

```
user@A> traceroute 192.168.0.4
traceroute to 192.168.0.4 (192.168.0.4), 30 hops max, 40 byte packets
 1  10.0.0.2 (10.0.0.2)  1.092 ms  1.034 ms  1.174 ms
 2  10.0.0.18 (10.0.0.18)  1.435 ms  2.062 ms  2.232 ms
    MPLS Label=299792 CoS=0 TTL=1 S=1
 3  bb04-cc1ab-1o0.spg1ab.juniper.net (192.168.0.4)  2.286 ms  1.432 ms  1.354 ms
```

Meaning The output shows that the LSP is used.

Related Documentation

- [Example: Enabling IS-IS Traffic Engineering Support on page 158](#)

Example: Enabling Wide IS-IS Metrics for Traffic Engineering

- [Understanding Wide IS-IS Metrics for Traffic Engineering on page 180](#)
- [Example: Enabling Wide IS-IS Metrics for Traffic Engineering on page 181](#)

Understanding Wide IS-IS Metrics for Traffic Engineering

All OSPF and IS-IS interfaces have a cost, which is a routing metric that is used in the link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. Unlike OSPF, in which the link metric is calculated automatically based on bandwidth, there is no automatic calculation for IS-IS. All IS-IS links use a metric of 10 by default.

Normally, IS-IS metrics can have values up to 63. The total cost to a destination is the sum of the metrics on all outgoing interfaces along a particular path from the source to the destination. By default, the total path metric is limited to 1023. This metric value is insufficient for large networks and provides too little granularity for traffic engineering, especially with high-bandwidth links. A wider range of metrics is also required if route leaking is used.

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 ($2^{24} - 1$).

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. To configure IS-IS to

generate only the new pair of TLVs and thus to allow the wider range of metric values, you must include the **wide-metrics-only** statement in the IS-IS configuration.

Example: Enabling Wide IS-IS Metrics for Traffic Engineering

This example shows how to allow a wide range of metric values on IS-IS interfaces.

- [Requirements on page 181](#)
- [Overview on page 181](#)
- [Configuration on page 181](#)
- [Verification on page 182](#)


Requirements

Before you begin, configure IS-IS on both routers. See “[Example: Configuring IS-IS](#)” on [page 14](#) for information about the sample IS-IS configuration.

Overview

[Figure 27 on page 181](#) shows the topology used in this example.

Figure 27: IS-IS Wide Metrics Topology



This example describes the steps on Device R1.

Configuration

CLI Quick Configuration	To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.
Device R1	<pre> set protocols isis level 2 wide-metrics-only set protocols isis level 1 wide-metrics-only set protocols isis interface lt-1/2/0.1 level 2 metric 100 set protocols isis interface lt-1/2/0.1 level 1 metric 100 </pre>
Step-by-Step Procedure	The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see <i>Using the CLI Editor in Configuration Mode</i> in the <i>CLI User Guide</i> .

To configure IS-IS checksums:

1. Configure a metric of 100 on the interface at both IS-IS levels.

```

[edit protocols isis interface lt-1/2/0.1]
user@R1# set level 2 metric 100
user@R1# set level 1 metric 100

```

2. Enable wide metrics.

```
[edit protocols isis]
user@R1# set level 2 wide-metrics-only
user@R1# set level 1 wide-metrics-only
```

Results From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols
isis {
  level 2 wide-metrics-only;
  level 1 wide-metrics-only;
  interface lt-1/2/0.1 {
    level 2 metric 100;
    level 1 metric 100;
  }
}
```

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

Verification

Confirm that the configuration is working properly.

Verifying That Wide Metrics Are Enabled

Purpose Make sure that the interface has the expected metric.

Action From operational mode, enter the **show isis interface extensive** command.

```
user@R1> show isis interface lt-1/2/0.1 extensive
IS-IS interface database:
lt-1/2/0.1
  Index: 68, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding
  Adjacency advertisement: Advertise
  Level 1
    Adjacencies: 1, Priority: 64, Metric: 100
    Hello Interval: 9.000 s, Hold Time: 27 s
    Designated Router: tp5-R2.02 (not us)
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 100
    Hello Interval: 9.000 s, Hold Time: 27 s
    Designated Router: tp5-R2.02 (not us)
```

Meaning The output shows that the metric is set to 100, as expected, at both Level 1 and Level 2.

Related Documentation

- [Example: Enabling IS-IS Traffic Engineering Support on page 158](#)
- [Example: Advertising Label-Switched Paths into IS-IS on page 172](#)

Example: Configuring Synchronization Between IS-IS and LDP

- [Understanding LDP-IGP Synchronization on page 183](#)
- [Example: Configuring Synchronization Between IS-IS and LDP on page 185](#)

Understanding LDP-IGP Synchronization

Synchronization between the Label Distribution Protocol (LDP) and the underlying interior gateway protocol (IGP) ensures that LDP is fully established before the IGP path is used for forwarding traffic.

LDP is often used to establish MPLS label-switched paths (LSPs) throughout a complete network domain using an IGP such as OSPF or IS-IS. In such a network, all links in the domain have IGP adjacencies as well as LDP adjacencies. LDP establishes the LSPs on the shortest path to a destination as determined by IP forwarding.

If the IGP and LDP are not synchronized, packet loss can occur. This issue is especially significant for applications such as a core network that does not employ BGP. Another example is an MPLS VPN where each provider edge (PE) router depends on the availability of a complete MPLS forwarding path to the other PE devices for each VPN that it serves. This means that along the shortest path between the PE routers, each link must have an operational hello adjacency and an operational LDP session, and MPLS label bindings must have been exchanged over each session.

LDP establishes MPLS LSPs along the shortest path to the destination as determined by IP forwarding. In a Layer 2 VPN or Layer 3 VPN scenario, if the LSP is not yet formed between the PE devices, services depending on MPLS forwarding fail. When LDP has not completed exchanging label bindings with an IGP next hop, traffic is discarded if the head end of the LSP forwards traffic because the LSP is assumed to be in place.

There are various reasons that the LSP fails to come up, as follows:

- Configuration errors and implementation issues.
- When an LDP hello adjacency or an LDP session with a peer is lost due to some error while the IGP still points to that peer. IP forwarding of traffic continues on the IGP link associated with the LDP peer rather than being shifted to another IGP link with which LDP is synchronized.
- When a new IGP link comes up, causing the next hop to a certain destination to change in the IGP's shortest-path-first (SPF) calculations. Although the IGP might be up on the new link, LDP might not have completed label exchange for all the routes. This condition might be transient or due to a misconfiguration.

LDP-IGP synchronization discourages a link from being used while the LDP sessions are not fully established. When LDP is not fully operational on a link, the IGP advertises a maximum cost for the link, thus preventing traffic from flowing through it. The IGP does not advertise the original cost or metric for the link until either LDP label exchange has been completed with the peer on the link or a configured amount of time has passed (the holddown period).

When synchronization is configured, LDP notifies the IGP to advertise the maximum cost for the link when one of the following triggering events takes place:

- The LDP hello adjacency goes down.
- The LDP session goes down.
- LDP is configured on an interface.

If the holddown timer has been configured, the timer starts when the triggering event takes place. When the timer expires, LDP notifies the IGP to resume advertising the original cost.

If the holddown timer has not been configured, the IGP waits (endlessly) until bindings have been received from downstream routers for all the forwarding equivalence classes (FECs) that have a next hop on that interface. Only after that takes place does LDP notify the IGP to bring down the cost on the interface.

LDP-IGP synchronization is supported only for directly connected peers and links with the platform label space.

Synchronization Behavior During Graceful Restart

LDP-IGP synchronization does not take place while the IGP is in the process of a graceful restart. When the graceful restart completes, links for which synchronization has been configured are advertised with maximum metrics in either of the following cases:

- LDP is not yet operational on the link and no holddown timer has been configured.
- The configured holddown timer has not expired.

During LDP graceful restart, no synchronization operations are done. If the LDP graceful restart is terminated, LDP notifies the IGPs to advertise the links with the maximum metric.

Synchronization Behavior on LAN Interfaces

LDP-IGP synchronization does not take place on LAN interfaces unless the IGP has a point-to-point connection over the LAN configured on the interface. The reason for this is that multiple LDP peers might be connected on such an interface unless a point-to-point connection to a single peer has been configured. Because synchronization raises the cost on the interface high enough to prevent traffic from being forwarded to that link, if multiple peers are connected, the cost is raised on all the peers even though LDP might be unsynchronized with only one of the peers. Consequently, traffic is diverted away from all the peers, an undesirable situation.

Synchronization Behavior on IGP Passive Interfaces

On IGP passive interfaces, the link cost is not raised when LDP-IGP synchronization is configured and a triggering event occurs.

Synchronization and TE Metrics

When traffic engineering is configured for an IGP, LDP-IGP synchronization does not affect the traffic engineering metric advertised for the link, regardless of whether the traffic-engineering (TE) metric is explicitly configured or the default value.

Example: Configuring Synchronization Between IS-IS and LDP

This example shows how to enable synchronization between IS-IS and LDP.

- [Requirements on page 185](#)
- [Overview on page 185](#)
- [Configuration on page 186](#)
- [Verification on page 187](#)

Requirements

Before you begin, configure IS-IS and LDP. For an example, see *Example: Configuring a Layer 3 VPN with Route Reflection and AS Override*.

Overview

LDP distributes labels in non-traffic-engineered applications. Labels are distributed along the best path determined by IS-IS. If the synchronization between LDP and IS-IS is lost, the label-switched path (LSP) goes down. Therefore, LDP and IS-IS synchronization is beneficial. When LDP synchronization is configured and when LDP is not fully operational on a given link (a session is not established and labels are not exchanged), IS-IS advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.

LDP synchronization is supported only on point-to-point interfaces and LAN interfaces configured as point-to-point interfaces under IS-IS. LDP synchronization is not supported during graceful restart.

To advertise the maximum cost metric until LDP is operational for LDP synchronization, include the **ldp-synchronization** statement:

```
ldp-synchronization {
  disable;
  hold-time seconds;
}
```


To disable synchronization, include the **disable** statement. To configure the time period to advertise the maximum cost metric for a link that is not fully operational, include the **hold-time** statement.



NOTE: When an interface has been in the holddown state for more than 3 minutes, a system log message with a warning level is sent. This message appears in both the messages file and the trace file.

Figure 28 on page 186 shows the topology used in this example.

Figure 28: IS-IS and LDP Synchronization Topology



This example describes the steps on Device P1.

Configuration

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

Device P1

```

set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis interface all
set protocols isis interface all ldp-synchronization
set protocols isis interface all point-to-point
set protocols isis interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable

```

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

To configure synchronization between IS-IS and LDP:

1. Enable MPLS on the interfaces


```

[edit protocols mpls]
user@P1# set interface all
user@P1# set interface fxp0.0 disable

```
2. Enable IS-IS on the interfaces.


```

[edit protocols isis]
user@P1# set interface all
user@P1# set interface fxp0.0 disable

```
3. Enable LDP on the interfaces.


```
[edit protocols ldp]
user@P1# set interface all
user@P1# set interface fxp0.0 disable
```

4. Enable LDP synchronization on the IS-IS interfaces.

```
[edit protocols isis interface all]
user@P1# set ldp-synchronization
```

5. Configure the IS-IS interfaces to behave like point-to-point interfaces.

```
[edit protocols isis interface all]
user@P1# set point-to-point
```

Results From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P1# show protocols
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
isis {
  interface all {
    ldp-synchronization;
    point-to-point;
  }
  interface fxp0.0 {
    disable;
  }
}
ldp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration on Device R2.

Verification

Confirm that the configuration is working properly.

Verifying LDP Synchronization

Purpose Check LDP synchronization setting on the IS-IS interfaces.

Action From operational mode, enter the **show isis interface extensive** command.

```
user@P1> show isis interface extensive
IS-IS interface database:
1o0.0
```

```
Index: 113, State: 0x6, Circuit id: 0x1, Circuit type: 0
LSP interval: 100 ms, CSNP interval: disabled, Loose Hello padding
Adjacency advertisement: Advertise
Level 1
  Adjacencies: 0, Priority: 64, Metric: 0
  Passive
Level 2
  Adjacencies: 0, Priority: 64, Metric: 0
  Passive
ge-1/2/0.0
  Index: 116, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 15 s, Loose Hello padding
  Adjacency advertisement: Advertise
  LDP sync state: in sync, for: 17:22:06, reason: LDP up during config
  config holddtime: infinity
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
ge-1/2/1.0
  Index: 114, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 15 s, Loose Hello padding
  Adjacency advertisement: Advertise
  LDP sync state: in sync, for: 17:22:06, reason: LDP up during config
  config holddtime: infinity
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
ge-1/2/2.0
  Index: 115, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 15 s, Loose Hello padding
  Adjacency advertisement: Advertise
  LDP sync state: in sync, for: 17:22:06, reason: LDP up during config
  config holddtime: infinity
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
```

Meaning The output shows that LDP is synchronized with IS-IS.

Related Documentation

- *Example: Configuring Synchronization Between LDP and OSPF*

CHAPTER 10

IS-IS Scaling and Throttling

- [Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces on page 189](#)
- [Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces on page 194](#)
- [Example: Configuring Mesh Groups of IS-IS Interfaces on page 199](#)

Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces

- [Understanding Link-State PDU Throttling for IS-IS Interfaces on page 189](#)
- [Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces on page 190](#)

Understanding Link-State PDU Throttling for IS-IS Interfaces

Link-state PDU throttling by use of the **lsp-interval** statement is a mechanism to control the flooding pace to neighboring routing devices to prevent overloading them.

Control traffic (link-state PDU and related packets) might cause delays in user traffic (information packets) because control traffic always has precedence in terms of scheduling on the interface cards.

Unfortunately, the control traffic transmission rate does not get lower on low-bandwidth interfaces such as DS-0 or fractional T1/E1 lines. Control traffic stays the same, regardless of line bandwidth.

Junos OS does not support automated calculation of link-state PDU throttling based on available bandwidth because the lowest-speed interface cards on a Juniper Networks routing device starts at T1/E1 speeds (1.5 and 2 Mbps). It is assumed that even with link-state PDU pacing of 20 ms, the control traffic will not consume more than half of the interface bandwidth.

However, there might be fractional T1/E1 circuits (less than the full bandwidth) configured as well, where link-state PDU pacing might have to be adjusted.

Thus, the **lsp-interval** statement helps to resolve two issues: regulating the control-traffic-to-user-traffic ratio, and protecting neighbors during transient situations.

The traffic subject to this pacing is non-self-originated traffic, which is traffic that has been originated by other routers, not the local router. Junos OS has hard-coded rate limiting for locally generated link-state PDUs. All the link-state PDUs are paced using a 20 ms timer. Additionally, there is logic that makes sure that the adjacency is reliably up for some time before advertising the adjacency.

Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces

This example shows how to modify the link-state PDU interval time.

- [Requirements on page 190](#)
- [Overview on page 190](#)
- [Configuration on page 191](#)
- [Verification on page 192](#)

Requirements

Before you begin, configure IS-IS. See “[Example: Configuring IS-IS](#)” on page 14 for information about the sample IS-IS configuration.

Overview

To keep reachability information in the network current, link-state protocols need to originate, distribute, and revoke or time-out topology information. In IS-IS, topology information is encoded in link-state PDUs.

By default, the routing device sends one link-state PDU out an interface every 100 milliseconds. To modify this interval, include the **lsp-interval** statement:

lsp-interval *milliseconds*;


To disable the transmission of all link-state PDUs, set the interval to 0.

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 and also to ensure that user traffic is not delayed on low-bandwidth links.

In this example, an IS-IS routing device on a LAN segment is configured to send link-state PDUs every 1000 milliseconds.

[Figure 29 on page 190](#) shows the topology used in this example.

Figure 29: IS-IS Link-State PDU Interval Topology



This example describes the steps on Device R1.

Configuration

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

Device R1

```
set interfaces fe-1/2/0 unit 0 description to-R2
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set protocols isis traceoptions file isis-trace
set protocols isis traceoptions flag lsp
set protocols isis interface fe-1/2/0.0 lsp-interval 1000
set protocols isis interface lo0.0
```

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

To configure the link-state PDU interval:

1. Configure the interfaces


```
[edit interfaces]
user@R1# set fe-1/2/0 unit 0 description to-R2
user@R1# set fe-1/2/0 unit 0 family inet address 10.0.0.1/30
user@R1# set fe-1/2/0 unit 0 family iso
user@R1# set lo0 unit 0 family inet address 192.168.0.1/32
user@R1# set lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
```
2. Enable IS-IS on the interfaces.


```
[edit protocols isis]
user@R1# set interface fe-1/2/0.0
user@R1# set interface lo0.0
```
3. Modify the link-state PDU interval.


```
[edit protocols isis interface fe-1/2/0.0]
user@R1# set lsp-interval 1000
```
4. (Optional) Enable tracing for tracking link-state PDU operations.


```
[edit protocols isis traceoptions]
user@R1# set file isis-trace
user@R1# set flag lsp
```

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

```
user@R1# show interfaces
fe-1/2/0 {
  unit 0 {
```

```
description to-R2;
family inet {
    address 10.0.0.1/30;
}
family iso;
}
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.1/32;
        }
        family iso {
            address 49.0002.0192.0168.0001.00;
        }
    }
}
}

user@R1# show protocols
isis {
    traceoptions {
        file isis-trace;
        flag lsp;
    }
    interface fe-1/2/0.0 {
        lsp-interval 1000;
    }
    interface lo0.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration on Device R2.

Verification

Confirm that the configuration is working properly.

- [Verifying the Link-State PDU Interval on page 192](#)
- [Checking the Link-State PDU Statistics on page 193](#)
- [Checking the Trace Log on page 193](#)

Verifying the Link-State PDU Interval

Purpose Check the link-state PDU interval setting on the IS-IS interface.

Action From operational mode, enter the **show isis interface extensive** command.

```
user@R1> show isis interface extensive
fe-1/2/0.0
Index: 70, State: 0x6, Circuit id: 0x1, Circuit type: 3
LSP interval: 1000 ms, CSNP interval: 10 s, Loose Hello padding
Adjacency advertisement: Advertise
Level 1
Adjacencies: 1, Priority: 64, Metric: 10
Hello Interval: 9.000 s, Hold Time: 27 s
Designated Router: R2.02 (not us)
```

```

Level 2
Adjacencies: 1, Priority: 64, Metric: 10
Hello Interval: 9.000 s, Hold Time: 27 s
Designated Router: R2.02 (not us)

```

Meaning The output shows that the link-state PDU interval is set to 1000 milliseconds.

Checking the Link-State PDU Statistics

Purpose Check the number of link-state PDUs sent and received.

Action From operational mode, enter the **show isis statistics** command.

```
user@R1> show isis statistics
```

IS-IS statistics for R1:

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	24	24	0	13	0
IIH	2467	24	0	836	0
CSNP	474	474	0	0	0
PSNP	0	0	0	0	0
Unknown	0	0	0	0	0
Totals	2965	522	0	849	0

Total packets received: 2965 Sent: 849

```

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0
SPF runs: 14
Fragments rebuilt: 15
LSP regenerations: 6
Purges initiated: 0

```

```
user@R2> show isis statistics
```

IS-IS statistics for R2:

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	13	13	0	24	0
IIH	828	15	0	2459	0
CSNP	0	0	0	474	0
PSNP	0	0	0	0	0
Unknown	0	0	0	0	0
Totals	841	28	0	2957	0

Total packets received: 841 Sent: 2957

```

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0
SPF runs: 17
Fragments rebuilt: 26
LSP regenerations: 11
Purges initiated: 0

```

Meaning The output shows the number of link-state PDUs sent and received on Device R1 and Device R2.

Checking the Trace Log

Purpose Check the IS-IS trace log to view the interval between packets.

Action From operational mode, enter the **show log isis-trace | match lsp** command.

```
user@R1> show log isis | match lsp
```

```
Jun 18 15:27:02.692031 Received L1 LSP R2.00-00, on interface fe-1/2/0.0
Jun 18 15:27:02.692753 Updating L1 LSP R2.00-00 in TED
Jun 18 15:27:44.396480 Updating L1 LSP R1.00-00 in TED
Jun 18 15:27:45.398077 Sending L1 LSP R1.00-00 on interface fe-1/2/0.0
Jun 18 15:28:44.689024 Received L1 LSP R2.02-00, on interface fe-1/2/0.0
Jun 18 15:28:44.689663 Updating L1 LSP R2.02-00 in TED
Jun 18 15:29:15.954900 Updating L2 LSP R1.00-00 in TED
Jun 18 15:29:16.955620 Sending L2 LSP R1.00-00 on interface fe-1/2/0.0
Jun 18 15:29:28.789986 Received L2 LSP R2.00-00, on interface fe-1/2/0.0
Jun 18 15:29:28.790620 Updating L2 LSP R2.00-00 in TED
Jun 18 15:30:27.727892 Received L2 LSP R2.02-00, on interface fe-1/2/0.0
Jun 18 15:30:27.728519 Updating L2 LSP R2.02-00 in TED
```

Meaning The output shows that Level 1 and Level 2 link-state PDUs are being sent and received roughly every 1000 milliseconds (1 second).

Related Documentation

- [Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces on page 194](#)

Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces

- [Understanding the Transmission Frequency for CSNPs on IS-IS Interfaces on page 194](#)
- [Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces on page 195](#)

Understanding the Transmission Frequency for CSNPs on IS-IS Interfaces

The complete sequence number PDU (CSNP) interval controls the frequency at which a routing device sends a directory of its link-state database.

When IS-IS is activated on a routing device's interface, the device first sends some IS-IS hello packets (IIHs) to its neighbors to ensure that the circuit is capable of transporting packets in both directions. In the IIHs, the router embeds information about the designated router (also called the designated intermediate system or DIS). One of the designated router roles on an IS-IS broadcast circuit is to synchronize the link-state databases on LANs. The designated router does this by periodically sending a directory of its link-state database, which is received by all the routing devices on a LAN.

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. The general recommendation is to use the default values or to increase the CSNP interval if there are a large number of broadcast circuits that need to be supplied with fresh CSNPs. Increasing the interval can help protect against CSNP flooding.

Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces

This example shows how to modify the complete sequence number PDU (CSNP) interval on IS-IS interfaces.

- [Requirements on page 195](#)
- [Overview on page 195](#)
- [Configuration on page 196](#)
- [Verification on page 197](#)

Requirements

Before you begin, configure IS-IS. See “[Example: Configuring IS-IS](#)” on page 14 for information about the sample IS-IS configuration.

Overview

CSNPs 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.

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. You might want to modify the default interval to protect against CSNP flooding.

To modify the CSNP interval, include the **csnp-interval** statement:

csnp-interval *seconds*;

The time can range from 1 through 65,535 seconds.


To configure the interface not to send any CSNPs, specify the **disable** option:

csnp-interval *disable*;

In this example, an IS-IS routing device on a LAN segment is configured to send CSNPs every 30 seconds.

[Figure 29 on page 190](#) shows the topology used in this example.

Figure 30: IS-IS CSNP Interval Topology



This example describes the steps on Device R1.

Configuration

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

Device R1

```
set interfaces fe-1/2/0 unit 0 description to-R2
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set protocols isis traceoptions file isis-trace
set protocols isis traceoptions flag csnp
set protocols isis interface fe-1/2/0.0 csnp-interval 30
set protocols isis interface lo0.0
```

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

To configure the CSNP interval:

1. Configure the interfaces


```
[edit interfaces]
user@R1# set fe-1/2/0 unit 0 description to-R2
user@R1# set fe-1/2/0 unit 0 family inet address 10.0.0.1/30
user@R1# set fe-1/2/0 unit 0 family iso
user@R1# set lo0 unit 0 family inet address 192.168.0.1/32
user@R1# set lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
```
2. Enable IS-IS on the interfaces.


```
[edit protocols isis]
user@R1# set interface fe-1/2/0.0
user@R1# set interface lo0.0
```
3. Modify the CSNP interval.


```
[edit protocols isis interface fe-1/2/0.0]
user@R1# set csnp-interval 30
```
4. (Optional) Enable tracing for tracking CSNP operations.


```
[edit protocols isis traceoptions]
user@R1# set file isis-trace
user@R1# set flag csnp
```

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

```
user@R1# show interfaces
fe-1/2/0 {
  unit 0 {
```

```

        description to-R2;
        family inet {
            address 10.0.0.1/30;
        }
        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.1/32;
        }
        family iso {
            address 49.0002.0192.0168.0001.00;
        }
    }
}
}

user@R1# show protocols
isis {
    traceoptions {
        file isis-trace;
        flag csnp;
    }
    interface fe-1/2/0.0 {
        csnp-interval 30;
    }
    interface lo0.0;
}

```

If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration on Device R2.

Verification

Confirm that the configuration is working properly.

- [Verifying the CSNP Interval on page 197](#)
- [Checking the CSNP Statistics on page 198](#)
- [Checking the IS-IS Log on page 199](#)

Verifying the CSNP Interval

Purpose Check the CSNP interval setting on the IS-IS interface.

Action From operational mode, enter the **show isis interface extensive** command.

```

user@R1> show isis interface extensive
IS-IS interface database:
fe-1/2/0.0
  Index: 70, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 30 s, Loose Hello padding
  Adjacency advertisement: Advertise
  Level 1
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s

```

```

    Designated Router: R2.02 (not us)
Level 2
Adjacencies: 1, Priority: 64, Metric: 10
Hello Interval: 9.000 s, Hold Time: 27 s
Designated Router: R2.02 (not us)

```

Meaning The output shows that the CSNP interval is set to 30 seconds.

Checking the CSNP Statistics

Purpose Checking the number of CSNPs sent and received.

Action From operational mode, enter the **show isis statistics** command.

```
user@R1> show isis statistics
```

```
IS-IS statistics for R1:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	5	5	0	3	0
IIH	94	20	0	43	0
CSNP	6	6	0	0	0
PSNP	0	0	0	0	0
Unknown	0	0	0	0	0
Totals	105	31	0	46	0

```
Total packets received: 105 Sent: 46
```

```

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0
SPF runs: 5
Fragments rebuilt: 5
LSP regenerations: 0
Purges initiated: 0
-----

```

```
user@R2> show isis statistics
```

```
IS-IS statistics for R2:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	3	3	0	5	0
IIH	35	11	0	86	0
CSNP	0	0	0	6	0
PSNP	0	0	0	0	0
Unknown	0	0	0	0	0
Totals	38	14	0	97	0

```
Total packets received: 38 Sent: 97
```

```

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0
SPF runs: 7
Fragments rebuilt: 7
LSP regenerations: 0
Purges initiated: 0

```

Meaning The output shows the number of CSNPs sent and received on Device R1 and Device R2.



NOTE: On broadcast links, only the designated intermediate system (DIS) sends CSNPs.

Checking the IS-IS Log

Purpose Check the IS-IS trace log to view the interval between packets.

Action From operational mode, enter the **show log isis-trace | match csn** command.

```
user@R1> show log isis-trace | match csn
```

```
Jun 18 14:36:19.504064 Received L1 CSN, source R2, interface fe-1/2/0.0
Jun 18 14:36:19.523065 Received L2 CSN, source R2, interface fe-1/2/0.0
Jun 18 14:36:48.904120 Received L1 CSN, source R2, interface fe-1/2/0.0
Jun 18 14:36:48.916425 Received L2 CSN, source R2, interface fe-1/2/0.0
Jun 18 14:37:14.954447 Received L1 CSN, source R2, interface fe-1/2/0.0
Jun 18 14:37:14.971329 Received L2 CSN, source R2, interface fe-1/2/0.0
Jun 18 14:37:44.227106 Received L1 CSN, source R2, interface fe-1/2/0.0
Jun 18 14:37:44.244181 Received L2 CSN, source R2, interface fe-1/2/0.0
```

Meaning The output shows that Level 1 and Level 2 CSNPs are being received roughly every 30 seconds.

Related Documentation

- [Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces on page 189](#)

Example: Configuring Mesh Groups of IS-IS Interfaces

- [Understanding IS-IS Mesh Groups on page 199](#)
- [Example: Configuring Mesh Groups of IS-IS Interfaces on page 200](#)

Understanding IS-IS Mesh Groups

A *mesh group* is a set of routing devices that are fully connected. That is, they have a fully meshed topology.

Junos OS supports IS-IS mesh groups as documented in RFC 2973, *IS-IS Mesh Groups*.

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.

Mesh groups provide a scaling method for the flooding subsystem. We recommend that you deploy mesh groups when your network design has a dense flooding topology. For example, consider the classical overlay topologies of the 1990s where 200 routers were fully meshed using permanent virtual circuits (PVCs) over an ATM core, because ATM was the only high-speed technology at the time. A PVC is a software-defined logical connection in a network such as a Frame Relay network.

What has changed since the 1990s is that IP and MPLS technology have reduced the ATM layer and removed the need for overlay meshing. The flooding graphs have become sparse in almost all practical deployments. In service provider networks, overlay networks are no longer used.

In enterprise networks, dense flooding graphs that, for example, lease a Layer 2 VPN service (an overlay network) to fully mesh its WAN routers might continue to be a useful architecture. In such cases, mesh groups might be useful.

Example: Configuring Mesh Groups of IS-IS Interfaces

This example shows how to configure mesh groups of IS-IS interfaces.

- [Requirements on page 200](#)
- [Overview on page 200](#)
- [Configuration on page 201](#)
- [Verification on page 204](#)

Requirements

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

Overview

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.

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:


mesh-group *value*;

To prevent an interface in the mesh group from flooding link-state PDUs, configure blocking on that interface:

mesh-group blocked;

Figure 31 on page 200 shows the topology used in this example.

Figure 31: IS-IS Mesh Topology



[“CLI Quick Configuration” on page 201](#) shows the configuration for all of the devices in [Figure 31 on page 200](#). The section [“Step-by-Step Procedure” on page 202](#) describes the steps on Device R1.

Configuration

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

Device R1

```
set interfaces fe-1/2/0 unit 0 description to-R2
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R4
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.5/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 description to-R3
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.9/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set protocols isis interface fe-1/2/0.0 mesh-group 1
set protocols isis interface fe-1/2/1.0 mesh-group 1
set protocols isis interface fe-1/2/2.0 mesh-group 1
set protocols isis interface lo0.0
```

Device R2

```
set interfaces fe-1/2/0 unit 0 description to-R1
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R1
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.14/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 description to-R1
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.21/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
set protocols isis interface fe-1/2/0.2 mesh-group 1
set protocols isis interface fe-1/2/1.0 mesh-group 1
set protocols isis interface fe-1/2/2.0 mesh-group 1
set protocols isis interface lo0.0
```

Device R3

```
set interfaces fe-1/2/0 unit 0 description to-R1
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.10/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R2
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.13/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 description to-R4
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0003.00
set protocols isis interface fe-1/2/0.0 mesh-group 1
```

```

set protocols isis interface fe-1/2/1.0 mesh-group 1
set protocols isis interface fe-1/2/2.0 mesh-group 1
set protocols isis interface lo0.0

```

Device R4

```

set interfaces fe-1/2/0 unit 0 description to-R1
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.6/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R3
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.18/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 description to-R2
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.22/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0004.00
set protocols isis interface fe-1/2/0.0 mesh-group 1
set protocols isis interface fe-1/2/1.0 mesh-group 1
set protocols isis interface fe-1/2/2.0 mesh-group 1
set protocols isis interface lo0.0

```

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

To configure an IS-IS mesh group:

1. Configure the interfaces.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 0 description to-R2
user@R1# set fe-1/2/0 unit 0 family inet address 10.0.0.1/30
user@R1# set fe-1/2/0 unit 0 family iso
user@R1# set fe-1/2/1 unit 0 description to-R4
user@R1# set fe-1/2/1 unit 0 family inet address 10.0.0.5/30
user@R1# set fe-1/2/1 unit 0 family iso
user@R1# set fe-1/2/2 unit 0 description to-R3
user@R1# set fe-1/2/2 unit 0 family inet address 10.0.0.9/30
user@R1# set fe-1/2/2 unit 0 family iso
user@R1# set lo0 unit 0 family inet address 192.168.0.1/32
user@R1# set lo0 unit 0 family iso address 49.0002.0192.0168.0001.00

```

2. Enable IS-IS on the interfaces, and assign a mesh group number.

```

[edit protocols isis]
user@R1# set interface fe-1/2/0.0 mesh-group 1
user@R1# set interface fe-1/2/1.0 mesh-group 1
user@R1# set interface fe-1/2/2.0 mesh-group 1
user@R1# set interface lo0.0

```

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

```

user@R1# show interfaces
fe-1/2/0 {
  unit 0 {

```



```

        description to-R2;
        family inet {
            address 10.0.0.1/30;
        }
        family iso;
    }
}
fe-1/2/1 {
    unit 0 {
        description to-R4;
        family inet {
            address 10.0.0.5/30;
        }
        family iso;
    }
}
fe-1/2/2 {
    unit 0 {
        description to-R3;
        family inet {
            address 10.0.0.9/30;
        }
        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.1/32;
        }
        family iso {
            address 49.0002.0192.0168.0001.00;
        }
    }
}
}

user@R1# show protocols
isis {
    interface fe-1/2/0.0 {
        mesh-group 1;
    }
    interface fe-1/2/1.0 {
        mesh-group 1;
    }
    interface fe-1/2/2.0 {
        mesh-group 1;
    }
    interface lo0.0;
}

```

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

Verification

Confirm that the configuration is working properly.

- [Checking the Interface Mesh Group on page 204](#)
- [Checking the IS-IS Statistics on page 205](#)

Checking the Interface Mesh Group

Purpose Verify that the mesh group is enabled on the IS-IS interfaces.

Action From operational mode, enter the **show isis interface extensive** command.

```
user@R1> show isis interface extensive
```

```
IS-IS interface database:
```

```
lo0.0
```

```
Index: 68, State: 0x6, Circuit id: 0x1, Circuit type: 0
LSP interval: 100 ms, CSNP interval: disabled, Loose Hello padding
Adjacency advertisement: Advertise
```

```
Level 1
```

```
Adjacencies: 0, Priority: 64, Metric: 0
Passive
```

```
Level 2
```

```
Adjacencies: 0, Priority: 64, Metric: 0
Passive
```

```
fe-1/2/0.1
```

```
Index: 73, State: 0x206, Circuit id: 0x1, Circuit type: 3
LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding
Adjacency advertisement: Advertise
CSNP interval: disabled, Mesh group: 1
```

```
Level 1
```

```
Adjacencies: 1, Priority: 64, Metric: 10
Hello Interval: 9.000 s, Hold Time: 27 s
Designated Router: tp5-R2.03 (not us)
```

```
Level 2
```

```
Adjacencies: 1, Priority: 64, Metric: 10
Hello Interval: 9.000 s, Hold Time: 27 s
Designated Router: tp5-R2.03 (not us)
```

```
fe-1/2/1.0
```

```
Index: 75, State: 0x206, Circuit id: 0x1, Circuit type: 3
LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding
Adjacency advertisement: Advertise
CSNP interval: disabled, Mesh group: 1
```

```
Level 1
```

```
Adjacencies: 1, Priority: 64, Metric: 10
Hello Interval: 9.000 s, Hold Time: 27 s
Designated Router: tp5-R4.03 (not us)
```

```
Level 2
```

```
Adjacencies: 1, Priority: 64, Metric: 10
Hello Interval: 9.000 s, Hold Time: 27 s
Designated Router: tp5-R4.03 (not us)
```

```
fe-1/2/2.0
```

```
Index: 76, State: 0x206, Circuit id: 0x1, Circuit type: 3
LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding
Adjacency advertisement: Advertise
CSNP interval: disabled, Mesh group: 1
```

```
Level 1
```

```
Adjacencies: 1, Priority: 64, Metric: 10
```

```

Hello Interval: 9.000 s, Hold Time: 27 s
Designated Router: tp5-R3.02 (not us)
Level 2
Adjacencies: 1, Priority: 64, Metric: 10
Hello Interval: 9.000 s, Hold Time: 27 s
Designated Router: tp5-R3.02 (not us)

```

Meaning Mesh group:1 in the output shows that the mesh group is enabled as expected.

Checking the IS-IS Statistics

Purpose Verify that the number of link-state PDUs received and sent is less than what it would be if the mesh group were not enabled.

Action From operational mode, enter the **show isis statistics** command.

```

user@R1> show isis statistics
IS-IS statistics for tp5-R1:
PDU type      Received    Processed    Drops      Sent      Reremit
LSP           73          73          0          37        0
IIH          4463         85          0         1525      0
CSNP         1294        1294         0           0         0
PSNP           0           0           0           2         0
Unknown        0           0           0           0         0
Totals       5830        1452         0         1564      0

```

Total packets received: 5830 Sent: 1564

```

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0
SPF runs: 26
Fragments rebuilt: 15
LSP regenerations: 4
Purges initiated: 0

```

Meaning After the adjacencies have been up for about 38 minutes, the output shows that Device R1 has received 73 link-state PDUs and sent 37 link-state PDUs. In the same topology in the same amount of time without the mesh group enabled, Device R1 would have received roughly 156 link-state PDUs and sent roughly 117 link-state PDUs.

Related Documentation

- [Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces on page 189](#)

CHAPTER 11

IS-IS CLNS

- [Example: Configuring IS-IS for CLNS on page 207](#)

Example: Configuring IS-IS for CLNS

- [Understanding IS-IS for CLNS on page 207](#)
- [Example: Configuring IS-IS for CLNS on page 207](#)

Understanding IS-IS for CLNS

IS-IS extensions provide the basic interior gateway protocol (IGP) support for collecting intradomain routing information for Connectionless Network Service (CLNS) destinations within a CLNS network. Routers that learn host addresses through End System-to-Intermediate System (ES-IS) can advertise the addresses to other routers (intermediate systems) by using IS-IS.

For more information about IS-IS, see the ISO 10589 standard.

Example: Configuring IS-IS for CLNS

This example shows how to create a routing instance and enable the IS-IS protocol on all interfaces.

- [Requirements on page 207](#)
- [Overview on page 207](#)
- [Configuration on page 208](#)
- [Verification on page 209](#)

Requirements

Before you begin, configure the network interfaces. See the *Junos OS Interfaces Configuration Guide for Security Devices*.

Overview

The configuration instructions in this topic describe how to create a routing instance called `aaaa`, enable IS-IS on all interfaces, define the BGP export policy name (`dist-bgp`), family (ISO), and protocol (BGP), and apply the export policy to IS-IS.

Configuration

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

```
set routing-instances aaaa protocols isis clns-routing
set routing-instances aaaa protocols isis interface all
set routing-instances aaaa protocols isis no-ipv4-routing no-ipv6-routing
set policy-options policy-statement dist-bgp from family iso protocol bgp
set policy-options policy-statement dist-bgp then accept
set routing-instances aaaa protocols isis export dist-bgp
```

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

To configure IS-IS for CLNS:

1. Enable CLNS routing.

```
[edit routing-instances aaaa]
user@host# set protocols isis clns-routing
```
2. Enable IS-IS on all interfaces.

```
[edit routing-instances aaaa]
user@host# set protocols isis interface all
```
3. (Optional) Disable IPv4 and IPv6 routing to configure a pure CLNS network.

```
[edit routing-instances aaaa]
user@host# set protocols isis no-ipv4-routing no-ipv6-routing
```
4. Define the BGP export policy name, family, and protocol.

```
[edit policy-options]
user@host# set policy-statement dist-bgp from family iso protocol bgp
```
5. Define the action for the export policy.

```
[edit policy-options]
user@host# set policy-statement dist-bgp then accept
```
6. Apply the export policy to IS-IS.

```
[edit routing-instances aaaa]
user@host# set protocols isis export dist-bgp
```

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

```
user@host# show routing-instances
aaaa {
  protocols {
    isis {
```

```
export dist-bgp;
no-ipv4-routing;
no-ipv6-routing;
clns-routing;
interface all;
}

user@host# show policy-options
policy-statement dist-bgp {
  from {
    family iso;
    protocol bgp;
  }
  then accept;
}
```

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

Verification

Confirm that the configuration is working properly.

- [Verifying the ISO Routes on page 209](#)
- [Checking the SPF Calculations on page 209](#)

Verifying the ISO Routes

Purpose Verify that the expected ISO routes are displayed in the IS-IS routing table.

Action From operational mode, enter the [show isis route](#) command.

Checking the SPF Calculations

Purpose Display information about IS-IS shortest-path-first (SPF) calculations.

Action From operational mode, enter the [show isis spf](#) command.

CHAPTER 12

Configuration Statements

- [\[edit protocols isis\] Hierarchy Level on page 213](#)
- [authentication-key \(Protocols IS-IS\) on page 216](#)
- [authentication-key-chain \(Protocols IS-IS\) on page 217](#)
- [authentication-type \(Protocols IS-IS\) on page 218](#)
- [backup-spf-options \(Protocols IS-IS\) on page 219](#)
- [bfd-liveness-detection \(Protocols IS-IS\) on page 220](#)
- [checksum \(Protocols IS-IS\) on page 222](#)
- [clns-routing on page 222](#)
- [clns-updown-compatibility on page 223](#)
- [context-identifier \(Protocols IS-IS\) on page 224](#)
- [csnp-interval on page 225](#)
- [disable \(Protocols IS-IS\) on page 226](#)
- [disable \(LDP Synchronization for IS-IS\) on page 227](#)
- [export \(Protocols IS-IS\) on page 228](#)
- [external-preference \(Protocols IS-IS\) on page 229](#)
- [family \(Protocols IS-IS\) on page 230](#)
- [graceful-restart \(Protocols IS-IS\) on page 231](#)
- [hello-authentication-key on page 232](#)
- [hello-authentication-key-chain on page 233](#)
- [hello-authentication-type on page 234](#)
- [hello-interval \(Protocols IS-IS\) on page 235](#)
- [hello-padding on page 236](#)
- [hold-time \(Protocols IS-IS\) on page 238](#)
- [hold-time \(LDP Synchronization for IS-IS\) on page 239](#)
- [ignore-attached-bit on page 240](#)
- [ignore-lsp-metrics \(Protocols IS-IS\) on page 241](#)
- [interface \(Protocols IS-IS\) on page 242](#)
- [ipv4-multicast on page 244](#)

- [ipv4-multicast-metric](#) on page 245
- [ipv6-multicast](#) on page 245
- [ipv6-multicast-metric](#) on page 246
- [ipv6-unicast](#) on page 247
- [ipv6-unicast-metric](#) on page 248
- [isis](#) on page 249
- [label-switched-path \(Protocols IS-IS\)](#) on page 250
- [ldp-synchronization](#) on page 251
- [level \(Global IS-IS\)](#) on page 252
- [level \(IS-IS Interfaces\)](#) on page 253
- [link-protection \(Protocols IS-IS\)](#) on page 254
- [loose-authentication-check](#) on page 254
- [lsp-equal-cost](#) on page 255
- [lsp-interval](#) on page 256
- [lsp-lifetime](#) on page 257
- [max-areas](#) on page 258
- [max-hello-size](#) on page 259
- [max-lsp-size](#) on page 259
- [max-snp-size](#) on page 260
- [mesh-group \(Protocols IS-IS\)](#) on page 261
- [metric \(Protocols IS-IS\)](#) on page 262
- [multicast-rpf-routes](#) on page 263
- [multipath \(Protocols IS-IS\)](#) on page 264
- [no-adjacency-down-notification \(Protocols IS-IS\)](#) on page 265
- [no-adjacency-holddown](#) on page 266
- [no-authentication-check](#) on page 266
- [no-csnp-authentication](#) on page 267
- [no-eligible-backup \(Protocols IS-IS\)](#) on page 267
- [no-hello-authentication](#) on page 268
- [no-ipv4-multicast](#) on page 268
- [no-ipv4-routing](#) on page 269
- [no-ipv6-multicast](#) on page 270
- [no-ipv6-routing](#) on page 271
- [no-ipv6-unicast](#) on page 272
- [no-psnp-authentication](#) on page 272
- [no-unicast-topology](#) on page 273
- [node-link-protection \(Protocols IS-IS\)](#) on page 273

- [overload \(Protocols IS-IS\)](#) on page 274
- [passive \(Protocols IS-IS\)](#) on page 277
- [point-to-point](#) on page 278
- [preference \(Protocols IS-IS\)](#) on page 279
- [prefix-export-limit \(Protocols IS-IS\)](#) on page 280
- [priority \(Protocols IS-IS\)](#) on page 281
- [protocols](#) on page 282
- [reference-bandwidth \(Protocols IS-IS\)](#) on page 284
- [rib-group \(Protocols IS-IS\)](#) on page 285
- [routing-instances \(Multiple Routing Entities\)](#) on page 286
- [shortcuts \(Protocols IS-IS\)](#) on page 287
- [spf-options \(Protocols IS-IS\)](#) on page 288
- [static-host-mapping](#) on page 290
- [te-metric \(Protocols IS-IS\)](#) on page 291
- [topologies \(Protocols IS-IS\)](#) on page 292
- [traceoptions \(Protocols IS-IS\)](#) on page 293
- [traffic-engineering \(Protocols IS-IS\)](#) on page 296
- [wide-metrics-only](#) on page 299

[\[edit protocols isis\]](#) Hierarchy Level

The following statement hierarchy can also be included at the [\[edit protocols isis\]](#) hierarchy level.

```

protocols {
  isis {
    disable;
    clns-routing;
    clns-updown-compatibility
    context-identifier ip-address</prefix> {
      level (1 | 2) <disable>;
    }
    export [ policy-names ];
    graceful-restart {
      disable;
      helper-disable;
      restart-duration seconds;
    }
    ignore-attached-bit;
    interface interface-name {
      ... the interface subhierarchy appears after the main [edit protocols isis] hierarchy ...
    }
    label-switched-path name level level metric metric;
    level (1 | 2) {
      disable;
      authentication-key key;
      authentication-type authentication;

```

```

    external-preference preference;
    no-csnp-authentication;
    no-hello-authentication;
    no-psnp-authentication;
    preference preference;
    prefix-export-limit number;
    wide-metrics-only;
}
loose-authentication-check;
lsp-lifetime seconds;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
no-ipv6-routing;
overload {
    advertise-high-metrics;
    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>;
}
traffic-engineering {
    disable;
    family inet {
        shortcuts {
            multicast-rpf-routes;
        }
    }
    family inet6 {
        shortcuts;
    }
}
ignore-lsp-metrics;
}

isis {
    interface interface-name {
        disable;
    }
}

```

```

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;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}
checksum;
csnp-interval (seconds | disable);
hello-padding (adaptive | loose | strict);
ldp-synchronization {
  disable;
  hold-time seconds;
}
level (1 | 2) {
  disable;
  hello-authentication-key key;
  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;
  te-metric metric;
}
link-protection;
lsp-interval milliseconds;
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;
}
}
}

```

- Related Documentation**
- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)
 - [\[edit protocols\] Hierarchy Level](#)

authentication-key (Protocols IS-IS)

Syntax	authentication-key <i>key</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	<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 (" ").



CAUTION: A simple password for authentication is truncated if it exceeds 254 characters.

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

- Related Documentation**
- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 45](#)

authentication-key-chain (Protocols IS-IS)

Syntax	authentication-key-chain <i>key-chain-name</i> ;
Hierarchy Level	[edit logical-systems <i>name</i> protocols isis level <i>level-number</i>], [edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols isis level <i>level-number</i>], [edit protocols isis level <i>level-number</i>], [edit routing-instances <i>instance-name</i> protocols isis level <i>level-number</i>]
Release Information	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Apply and enable an authentication keychain to the routing device.
Options	key-chain —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	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 on page 46 • <i>Example: Configuring Route Authentication for BGP</i> • <i>Example: Configuring BFD Authentication for Static Routes</i> • <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i> • Understanding Hitless Authentication Key Rollover for IS-IS on page 45

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 on page 45• authentication-key on page 216• no-authentication-check on page 266

backup-spf-options (Protocols IS-IS)

Syntax	<pre>backup-spf options { per-prefix-calculation; }</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	Statement introduced in Junos OS Release 11.4R3.
Description	<p>Configure options for running the shortest-path-first (SPF) algorithm for backup next hops for protected interfaces. Use these options to override the default behavior of having Junos OS calculate backup paths for all the topologies in an instance when at least one interface is configured with link protection or node-link protection.</p> <p>This configuration option is used in conjunction with the node-link-protection statement to compute per-prefix loop-free alternate routes (LFAs).</p>
Options	per-prefix-calculation —Calculate backup nexthops for non-best prefix originators.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control-level—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring MPLS Egress Protection for Layer 3 VPN Services</i> • <i>Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP</i> • <i>link-protection</i> • <i>node-link-protection</i>

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 interface-name],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name],</p> <p>[edit protocols isis interface interface-name],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]</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>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 statement introduced in Junos OS Release 9.0.</p> <p>authentication algorithm, authentication key-chain, and authentication loose-check options introduced in Junos OS Release 9.6.</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, 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 on page 91 • Example: Configuring BFD Authentication for IS-IS on page 99
------------------------------	---


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 on page 50

clns-routing

Syntax	clns-routing;
Hierarchy Level	[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 11.2 for MX Series routers.
Description	Enable IS-IS to exchange Connectionless Network Service (CLNS) routes. CLNS is a Layer 3 protocol, similar to IPv4. CLNS uses network service access points (NSAPs) to address end systems and intermediate systems.
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 for CLNS on page 207• Example: Configuring Static Routes for CLNS• Example: Configuring BGP and CLNS

clns-updown-compatibility

Syntax	clns-updown-compatibility;
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	Statement introduced in Junos OS Release 11.4R7.
Description	<p>When you enable IS-IS to exchange Connectionless Network Service (CLNS) routes, Junos OS sets the reserved (R) bit in the default metric field inside type, length, and value (TLV) type-3 (ES-Neighbor) as a marker for routing loop prevention. Junos OS uses the up/down bit for marking prefixes on the Level 2-to-Level 1 boundary as being propagated Down, such that any router in that area never propagates it Up on a Level 1-to-Level 2 boundary. For detailed information about how this works in IP routing environments, see RFC 2966, <i>Domain-wide Prefix Distribution with Two-Level IS-IS</i>.</p> <p>Some other vendors' platforms might not support up/down bit setting in CLNS route TLVs. If one of these vendors' platforms receives this TLV with the R bit set, the platform discards the information.</p> <p>When you use the clns-updown-compatibility statement in the IS-IS configuration, the R bit is set to 0, and the issue is resolved. The clns-updown-compatibility statement causes Junos OS to use the Internal/External metric-type bit in the TLV header instead of using the R bit as the up/down bit marker. This has the advantage that older end system (ES) equipment does not receive TLV headers with the R bit set.</p>
<div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p>CAUTION: Not using the R bit can lead to potential routing loops. You can use the site-of-origin (SoO) extended community to prevent a looped BGP update from being injected back to IS-IS when received from a remote provider edge (PE) device. The receiving PE device can check against the SoO community, and if the value matches its own, the NLRI is not accepted.</p> </div> </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 for CLNS on page 207 • Example: Configuring Static Routes for CLNS • Example: Configuring BGP and CLNS

context-identifier (Protocols IS-IS)

Syntax	context-identifier <i>identifier</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit protocols isis]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Configure IS-IS context identifier information.
Options	<i>identifier</i> —IPv4 address that defines a protection pair. The context identifier is manually configured on both primary and protector PEs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show isis context-identifier on page 349

csnp-interval

Syntax	<code>csnp-interval (seconds disable);</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 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 multiplied by the number of IS-IS adjacencies over point-to-point links, which are in UP state.</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 multiplied by the number of IS-IS adjacencies over point-to-point links, which are in UP state.
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	<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 CSNP Packets on IS-IS Interfaces on page 195


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 on page 20 • IS-IS Overview on page 3

disable (LDP Synchronization for IS-IS)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name ldp-synchronization], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name ldp-synchronization], [edit protocols isis interface interface-name ldp-synchronization], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name ldp-synchronization]
Release Information	Statement introduced in Junos OS Release 7.5.
Description	Disable LDP synchronization for IS-IS.
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 Synchronization Between IS-IS and LDP on page 183

export (Protocols IS-IS)

Syntax	<code>export [<i>policy-names</i>];</code>
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	<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> 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.</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 on page 61• Example: Configuring an IS-IS Default Route Policy on Logical Systems on page 82

external-preference (Protocols IS-IS)

Syntax	<code>external-preference <i>preference</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis level level-number],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number],</p> <p>[edit protocols isis level level-number],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]</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 preference of external routes.
Options	<p><i>preference</i>—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	<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 • Example: Redistributing OSPF Routes into IS-IS on page 61 • Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS on page 72 • preference on page 279

family (Protocols IS-IS)

Syntax	<pre>family inet { shortcuts { multicast-rpf-routes; } } family inet6 { shortcuts; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering], [edit protocols isis traffic-engineering], [edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering]
Release Information	Statement introduced in Junos OS Release 9.3. Support for IPv6 for IGP shortcuts introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the address family for traffic engineering IS-IS interior gateway protocol (IGP) shortcuts.
Options	inet—IPv4 address family inet6—IPv6 address family 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">•

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"> • <i>Configuring Routing Protocols Graceful Restart</i>


hello-authentication-key

Syntax	<code>hello-authentication-key password;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level number], [edit protocols isis interface <i>interface-name</i> level number], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> 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 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 216• authentication-type on page 218• hello-authentication-type on page 234

hello-authentication-key-chain

Syntax	hello-authentication-key-chain <i>key-chain-name</i> ;
Hierarchy Level	[edit logical-systems <i>name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>name</i> routing-instances <i>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>instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Apply an authentication keychain to the IS-IS interface.
Options	<i>key-chain-name</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	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 on page 45

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 number],</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 number],</p> <p>[edit protocols isis interface <i>interface-name</i> level number],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level number]</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 style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <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> </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 216 • authentication-type on page 218 • hello-authentication-key on page 232

hello-interval (Protocols IS-IS)

Syntax	<code>hello-interval <i>seconds</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> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
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	<p>seconds—Frequency of transmission for hello packets.</p> <p>Range: 1 through 20,000 seconds</p> <p>Default: 3 seconds (for designated intermediate system [DIS] routers), 9 seconds (for non-DIS routers)</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"> • hold-time on page 239

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 on page 13](#)

hold-time (Protocols IS-IS)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code> <code>[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>],</code> <code>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</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>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 on page 13• hello-interval on page 235

hold-time (LDP Synchronization for 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> ldp-synchronization],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ldp-synchronization],</p> <p>[edit protocols isis interface <i>interface-name</i> ldp-synchronization],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ldp-synchronization]</p>
Release Information	Statement introduced in Junos OS Release 7.5.
Description	Configure the time period to advertise the maximum cost metric for a link that is not fully operational.
<div>  <p>NOTE: When an interface has been in the holddown state for more than 3 minutes, a system log message with a warning level is sent. This message appears in both the messages file and the trace file.</p> </div>	
Options	<p>seconds—Hold-time value, in seconds.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: Infinity</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 Synchronization Between IS-IS and LDP on page 183

ignore-attached-bit

Syntax	ignore-attached-bit;
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	<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	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">•

ignore-lsp-metrics (Protocols IS-IS)


Syntax	ignore-lsp-metrics;
Hierarchy Level	<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 traffic-engineering],</p> <p>[edit protocols isis traffic-engineering],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering]</p>
Release Information	Statement introduced in Junos OS Release 8.0.
Description	<p>Ignore the metrics for RSVP label-switched paths (LSPs) in IS-IS traffic engineering shortcut calculations or when you configure LDP over RSVP LSPs.</p> <p>If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops. Ignoring the metric of RSVP LSPs avoids mutual dependency between IS-IS and RSVP, eliminating the time period when the RSVP metric used for tunneling traffic is not up to date.</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 Synchronization Between IS-IS and LDP on page 183 • Example: Enabling IS-IS Traffic Engineering Support on page 158 • Example: Enabling Wide IS-IS Metrics for Traffic Engineering on page 180 • Example: Advertising Label-Switched Paths into IS-IS on page 172 • shortcuts on page 287

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	[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	<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	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 on page 13• Example: Configuring Multi-Level IS-IS on page 19

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> 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.</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 on page 109

ipv4-multicast-metric

Syntax	<code>ipv4-multicast-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 multicast topology metric value for the level.
Options	metric —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 on page 109


ipv6-multicast

Syntax	<code>ipv6-multicast;</code>
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 on page 109

ipv6-multicast-metric

Syntax	ipv6-multicast-metric <i>metric</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level level-number], [edit protocols isis interface <i>interface-name</i> level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> 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.
Description	Specify the IPv6 alternate 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 on page 109

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	Configure alternate IPv6 unicast topologies. 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.
<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 on page 131

ipv6-unicast-metric

Syntax	<code>ipv6-unicast-metric <i>metric</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level level-number], [edit protocols isis interface <i>interface-name</i> level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> 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.
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	<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 IPv4 and IPv6 Unicast Topologies on page 131

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 on page 14 • Example: Configuring Multi-Level IS-IS on page 20

label-switched-path (Protocols IS-IS)

Syntax	label-switched-path <i>name</i> level <i>level-number</i> metric <i>metric</i> ;
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.
Description	<p>Advertise label-switched paths (LSPs) into IS-IS as point-to-point links. The LSP is advertised in the appropriate IS-IS levels as a point-to-point link and contains a local address and a remote address.</p> <p>When you advertise LSPs into IS-IS as point-to-point links, the LSPs are used in SPF calculations. The advertisement contains a local address (the from address of the LSP), a remote address (the to address of the LSP), and a metric.</p> <p>Before a single-hop LSP between a multiaccess link can be announced as up and used in SPF calculations, you must configure an LSP in both directions between two label-switched routers.</p>
Options	<p><i>name</i>—Identifies the LSP.</p> <p><i>level-number</i>—IS-IS level number.</p> <p>Values: 1 or 2</p> <p><i>metric</i>—Metric value.</p> <p>Range: 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)</p> <p>Default: 0 (for lo0), 10 (for all other interfaces)</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: Advertising Label-Switched Paths into IS-IS on page 172

ldp-synchronization

Syntax	<pre>ldp-synchronization { disable; hold-time seconds; }</pre>
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	Statement introduced in Junos OS Release 7.5.
Description	<p>Enable synchronization by advertising the maximum cost metric until LDP is operational on the link.</p> <p>LDP distributes labels in non-traffic-engineered applications. Labels are distributed along the best path determined by IS-IS. If the synchronization between LDP and IS-IS is lost, the label-switched path (LSP) goes down. Therefore, IS-IS and LDP synchronization is beneficial. When LDP synchronization is configured and when LDP is not fully operational on a given link (a session is not established and labels are not exchanged), IS-IS advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.</p> <p>LDP synchronization is supported only on point-to-point interfaces and LAN interfaces configured as point-to-point interfaces under IS-IS. LDP synchronization is not supported during graceful restart.</p> <p>To advertise the maximum cost metric until LDP is operational for LDP synchronization, include the ldp-synchronization statement.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Synchronization Between LDP and OSPF</i>

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>; disable; external-preference <i>preference</i>; no-csnp-authentication; no-hello-authentication; no-psnp-authentication; preference <i>preference</i>; wide-metrics-only; }</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	<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>level-number—IS-IS level number.</p> <p>Values: 1 or 2</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 IS-IS on page 13](#)
 - [Example: Configuring Multi-Level IS-IS on page 19](#)

level (IS-IS Interfaces)

Syntax	<pre> level <i>level-number</i> { level (IS-IS Interfaces); hello-authentication-key <i>key</i>; hello-authentication-key-chain <i>key-chain-name</i>; hello-authentication-type <i>authentication</i>; hello-interval <i>seconds</i>; hold-time <i>seconds</i>; ipv4-multicast-metric <i>number</i>; ipv6-unicast-metric <i>number</i>; metric <i>metric</i>; passive; priority <i>number</i>; te-metric <i>metric</i>; } </pre>
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>
Description	<p>Configure the IS-IS level. You can configure one instance of Level 1 routing and one instance of Level 2 routing 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>Default: The routing device operates as both a Level 1 and Level 2 router.</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 on page 13 • Example: Configuring Multi-Level IS-IS on page 19

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 on page 141• node-link-protection on page 273

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 on page 45

lsp-equal-cost

Syntax	<code>lsp-equal-cost;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering multipath], [edit protocols isis traffic-engineering multipath]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	<p>Configure label-switched paths (LSPs) to be retained as equal cost paths for load balancing when a better path metric is found during the IS-IS internal routing table calculation.</p> <p>When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) combinations can be lost during route calculation. To override this default IS-IS behavior, include the lsp-equal-cost statement for load balancing, so that the equal cost path information is retained in the 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">• Example: Enabling IS-IS Traffic Engineering Support on page 158• multipath on page 264• traffic-engineering on page 296

lsp-interval

Syntax	<code>lsp-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name],</p> <p>[edit protocols isis interface interface-name],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]</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>milliseconds—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 on page 189

lsp-lifetime

Syntax	<code>lsp-lifetime <i>seconds</i>;</code>
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>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 routing device 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 show isis overview 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	<p><i>seconds</i>—link-state PDU lifetime, in seconds.</p> <p>Range: 350 through 65,535 seconds</p> <p>Default: 1200 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 the Transmission Frequency for Link-State PDUs on IS-IS Interfaces on page 189 • http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

max-areas

Syntax	<code>max-areas <i>number</i>;</code>
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 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<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	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 Multi-Level IS-IS on page 19

max-hello-size

Syntax	<code>max-hello-size size;</code>
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 12.1.
Description	Modify the maximum size of IS-IS hello packets. IS-IS sends hello packets out of all IS-IS enabled interfaces to discover neighbors and form adjacencies between the devices.
Options	size —Maximum size allocated for IS-IS hello packets. Range: 512 through 1492 bytes Default: 1492 bytes
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 on page 13

max-lsp-size

Syntax	<code>max-lsp-size size;</code>
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 12.1.
Description	Modify the maximum size of IS-IS link-state PDUs. IS-IS sends link-state PDUs out of IS-IS enabled interfaces to distribute routing information between the IS-IS nodes.
Options	size —Maximum size allocated for IS-IS link-state PDUs. Range: 512 through 1492 bytes Default: 1492 bytes
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 on page 13

max-snp-size

Syntax	max-snp-size <i>size</i> ;
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 12.1.
Description	Modify the maximum size of partial or complete IS-IS sequence number PDUs. IS-IS sends sequence number packets out of IS-IS enabled interfaces to control the distribution of link-state PDUs between the IS-IS nodes. Sequence number packets provide a mechanism to synchronize the link-state databases of routers in the same area.
Options	size —Maximum size allocated for sequence number of partial or complete IS-IS packets. Range: 512 through 1400 bytes Default: 1400 bytes
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 on page 13

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 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 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 on page 199

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 7 on page 262 depicts IS-IS route export metric default values.

Table 7: 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 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 Engineering on page 180 • te-metric on page 291 • wide-metrics-only on page 299

multicast-rpf-routes

Syntax	multicast-rpf-routes;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering family inet shortcuts],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances traffic-engineering family inet shortcuts],</p> <p>[edit protocols isis traffic-engineering family inet shortcuts],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering family inet shortcuts]</p>
Release Information	Statement introduced in Junos OS Release 9.3.
Description	<p>Install unicast IPv4 routes into the multicast routing table (inet.2) for multicast reverse-path-forwarding (RPF) checks.</p> <p>Traffic engineering shortcuts must be enabled. IPv4 multicast topology must not be enabled. Label-switched paths (LSPs) must not be advertised into IS-IS.</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: Enabling IS-IS Traffic Engineering Support on page 158

multipath (Protocols IS-IS)

Syntax	<pre>multipath { lsp-equal-cost; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering], [edit protocols isis traffic-engineering]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	Enable load balancing for multiple label-switched paths (LSPs).
Options	<p>lsp-equal-cost—Configure LSPs to be retained as equal cost paths for load balancing when a better route metric is added to the routing table.</p> <p>When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) path combinations can be lost during route calculation. To override this default IS-IS behavior, include the lsp-equal-cost statement for load balancing, so that the equal cost path information is retained in the 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">• Example: Enabling IS-IS Traffic Engineering Support on page 158• lsp-equal-cost on page 255• <i>multipath (Protocols BGP)</i>• <i>multipath (Routing Options)</i>• traffic-engineering on page 296

no-adjacency-down-notification (Protocols IS-IS)

Syntax	no-adjacency-down-notification;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.0.
Description	<p>Disable adjacency down notification for IS-IS to allow for migration from IS-IS to OSPF without disruption of the RSVP neighbors and associated RSVP-signaled label-switched paths (LSPs).</p> <p>Whenever IS-IS is deactivated, the IS-IS adjacencies are brought down. IS-IS signals to RSVP to bring down any RSVP neighbors associated with the IS-IS adjacencies, and this further causes the associated LSPs signaled by RSVP to go down as well.</p> <p>A similar process occurs whenever OSPF is deactivated. The OSPF neighbors are brought down. OSPF signals to RSVP to bring down any of the RSVP neighbors associated with the OSPF neighbors, and this further causes the associated LSPs signaled by RSVP to go down as well.</p> <p>If you need to migrate from IS-IS to OSPF or from OSPF to IS-IS, the internal gateway protocol (IGP) notification to RSVP for an adjacency or neighbor down event needs to be ignored. Using the no-adjacency-down-notification or no-neighbor-down-notification statements, you can disable IS-IS adjacency down notification or OSPF neighbor down notification, respectively, until the migration is complete. The network administrator is responsible for configuring the statements before the migration, and then removing them from the configuration afterward, so that IGP notification can function normally.</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"> <i>no-neighbor-down-notification</i>

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 238

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">•• csnp-interval on page 225• hello-authentication-type on page 234

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 225

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 on page 141 • link-protection on page 254 • node-link-protection on page 273


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 234

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 on page 109

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.
	<p> NOTE: Note: Even when no-ipv4-routing 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>
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 on page 131

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 on page 109

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 on page 131

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 on page 131

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 on page 43

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 on page 109

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 on page 141 • link-protection on page 254

overload (Protocols IS-IS)

Syntax	<pre>overload { advertise-high-metrics; allow-route-leaking; timeout <i>seconds</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 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 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.

The **advertise-high-metric** setting is only valid while the routing device is in overload mode. When **advertise-high-metric** is configured, IS-IS does not set the overload bit. Rather, it sets the metric to 63 or 16,777,214, depending whether wide metrics are enabled. This allows the overloaded routing device to be used for transit as a last resort.

An L1-L2 router in overload mode stops leaking route information between L1 and L2 levels and clears its attached bit. This is also true when **advertise-high-metrics** is configured.

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 on page 13](#)

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>],</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> 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 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 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>
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 on page 20 • disable on page 227

point-to-point

Syntax	point-to-point;
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 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 on page 3• Understanding IS-IS Designated Routers on page 39• Example: Configuring Synchronization Between IS-IS and LDP on page 183

preference (Protocols IS-IS)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>],</p> <p>[edit protocols isis level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis 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>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	<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 • Example: Redistributing OSPF Routes into IS-IS on page 61 • Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS on page 72 • external-preference on page 229

prefix-export-limit (Protocols IS-IS)

Syntax	<code>prefix-export-limit <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis level level-number],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number],</p> <p>[edit protocols isis level level-number],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]</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 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 on page 72 • Example: Redistributing OSPF Routes into IS-IS on page 61

priority (Protocols IS-IS)

Syntax	<code>priority <i>number</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> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
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	<p><i>number</i>—Priority value.</p> <p>Range: 0 through 127</p> <p>Default: 64</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 Designated Routers on page 39

protocols

```

Syntax protocols {
    bgp {
        ... bgp-configuration ...
    }
    isis {
        ... isis-configuration ...
    }
    ldp {
        ... ldp-configuration ...
    }
    msdp {
        ... msdp-configuration ...
    }
    mstp {
        ... mstp-configuration ...
    }
    ospf {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (iana | vendor);
        ... ospf-configuration ...
    }
    ospf3 {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (iana | vendor);
        ... ospf3-configuration ...
    }
    pim {
        ... pim-configuration ...
    }
    rip {
        ... rip-configuration ...
    }
    ripng {
        ... ripng-configuration ...
    }
    rstp {
        rstp-configuration;
    }
    vstp {
        vstp configuration;
    }
    vpls {
        vpls configuration;
    }
}

```

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

Release Information Statement introduced before Junos OS Release 7.4.

Support for RIPv6 introduced in Junos OS Release 9.0.
 Statement introduced in Junos OS Release 11.1 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Specify the protocol for a routing instance. You can configure multiple instances of many protocol types. Not all protocols are supported on the switches. See the switch CLI.

Options **bgp**—Specify BGP as the protocol for a routing instance.
isis—Specify IS-IS as the protocol for a routing instance.
ldp—Specify LDP as the protocol for a routing instance.
l2vpn—Specify Layer 2 VPN as the protocol for a routing instance.
msdp—Specify the Multicast Source Discovery Protocol (MSDP) for a routing instance.
mstp—Specify the Multiple Spanning Tree Protocol (MSTP) for a virtual switch routing instance.
ospf—Specify OSPF as the protocol for a routing instance.
ospf3—Specify OSPF version 3 (OSPFv3) as the protocol for a routing instance.



NOTE: OSPFv3 supports the **no-forwarding**, **virtual-router**, and **vrf** routing instance types only.

pim—Specify the Protocol Independent Multicast (PIM) protocol for a routing instance.
rip—Specify RIP as the protocol for a routing instance.
ripng—Specify RIP next generation (RIPv6) as the protocol for a routing instance.
rstp—Specify the Rapid Spanning Tree Protocol (RSTP) for a virtual switch routing instance.
vstp—Specify the VLAN Spanning Tree Protocol (VSTP) for a virtual switch routing instance.
vpls—Specify VPLS as the protocol for a routing instance.

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

Related Documentation

- *Example: Configuring Multiple Routing Instances of OSPF*

reference-bandwidth (Protocols IS-IS)

Syntax	<code>reference-bandwidth <i>reference-bandwidth</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 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>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	<i>reference-bandwidth</i> —Reference bandwidth value in bits per second. Range: 9600 through 1,000,000,000,000 bps 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: Configuring IS-IS on page 13• http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

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"> • <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i> • <i>Example: Importing Direct and Static Routes Into a Routing Instance</i> • <i>Understanding Multiprotocol BGP</i>

routing-instances (Multiple Routing Entities)

Syntax	<code>routing-instances <i>routing-instance-name</i> { ... }</code>
Hierarchy Level	[edit], [edit logical-systems <i>logical-system-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Configure an additional routing entity for a router. You can create multiple instances of BGP, IS-IS, OSPF, OSPFv3, and RIP for a router. You can also create multiple routing instances for separating routing tables, routing policies, and interfaces for individual wholesale subscribers (retailers) in a Layer 3 wholesale network.</p> <p>Each routing instance consist of the following:</p> <ul style="list-style-type: none">• A set of routing tables• A set of interfaces that belong to these routing tables• A set of routing option configurations <p>Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name my-instance, its corresponding IP unicast table is my-instance.inet.0. All routes for my-instance are installed into my-instance.inet.0.</p> <p>Routes are installed into the default routing instance inet.0 by default, unless a routing instance is specified.</p> <p>In Junos OS Release 9.0 and later, you can no longer specify a routing-instance name of <i>master</i>, <i>default</i>, or <i>bgp</i> or include special characters within the name of a routing instance.</p> <p>In Junos OS Release 9.6 and later, you can include a slash (/) in a routing-instance name only if a logical system is not configured. That is, you cannot include the slash character in a routing-instance name if a logical system other than the default is explicitly configured. Routing-instance names, further, are restricted from having the form <code>__.*__</code> (beginning and ending with underscores). The colon : character cannot be used when multipotology routing (MTR) is enabled.</p>
Default	Routing instances are disabled for the router.
Options	<p><i>routing-instance-name</i>—Name of the routing instance. This must be a non-reserved string of not more than 128 characters.</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 Interprovider Layer 3 VPN Option A](#)
 - [Example: Configuring Interprovider Layer 3 VPN Option B](#)
 - [Example: Configuring Interprovider Layer 3 VPN Option C](#)

shortcuts (Protocols IS-IS)

Syntax	<pre>shortcuts { multicast-rpf-routes; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering family (inet inet6)],</p> <p>[edit protocols isis traffic-engineering family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering family (inet inet6)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>The family statement and support for IPv6 routes for IS-IS traffic engineering shortcuts introduced in Junos OS Release 9.3.</p>
Description	<p>Configure IS-IS to use MPLS label-switched paths (LSPs) as next hops if possible when installing routing information into the inet.3 or inet6.3 routing table. Internal gateway protocol (IGP) shortcuts allow the IGP to install prefixes in inet.3 or inet6.3.</p> <p>It is only necessary to enable IGP shortcuts on the ingress router because that is the router performing the shortest-path-first (SPF) calculations.</p> <p>It is important to understand how IGP shortcuts affect the protocol and routing table relationship. The IGP performs SPF calculations to subnets downstream of LSP egress points, but the results of these calculations are entered into the inet.3 table only. At the same time, the IGP performs its traditional SPF calculations and enters the results of these calculations into the inet.0 table. The result is that although the IGP is making entries into the inet.3 table, BGP is still the only protocol with visibility into that table for the purposes of route resolution. Therefore, forwarding to AS-internal destinations still uses the inet.0 IGP routes, and the LSPs are only used for BGP next-hop resolution. If you want the LSPs to be used for IGP next-hop resolution, you must configure traffic-engineering bgp-igp.</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"> • Example: Enabling IS-IS Traffic Engineering Support on page 158 • traffic-engineering (Protocols IS-IS) on page 296 • traffic-engineering (Protocols MPLS)

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

- [Example: Configuring Link and Node Protection for IS-IS Routes on page 141](#)

static-host-mapping

Syntax	<pre>static-host-mapping { hostname { alias [<i>aliases</i>]; inet [<i>addresses</i>]; sysid <i>system-identifier</i>; } }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	(Optional) Statically map a hostname to one or more IP addresses and aliases, and configure an International Organization for Standardization (ISO) system identifier (system ID).
Default	If you do not statically map the hostname, the mapping is generated dynamically, based on the system configuration. For instance, if you omit the static-host-mapping hostname sysid statement, the IS-IS system ID is dynamically generated from the host portion of the ISO address configured on the loopback interface (lo0) and is mapped to the host-name statement configured at the [edit system] hierarchy level.
Options	<p>alias <i>alias</i>—Alias for the hostname.</p> <p>hostname—Fully qualified hostname.</p> <p>inet <i>address</i>—IP address. You can specify one or more IP addresses for the host.</p> <p>sysid <i>system-identifier</i>—ISO system identifier (system ID). This is the 6-byte portion of the Intermediate System-to-Intermediate System (IS-IS) network service access point (NSAP). We recommend that you use the host's IP address represented in binary-coded decimal (BCD) format. For example, the IP address 208.197.169.18 is 2081.9716.9018 in BCD.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Hostname of the Router or Switch</i>

te-metric (Protocols IS-IS)

Syntax	<code>te-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	Statement introduced before Junos OS Release 7.4.
Description	<p>Set the metric value used by traffic engineering for information injected into the traffic engineering database. The value of the traffic engineering metric does not affect normal IS-IS forwarding.</p> <p>When traffic engineering is enabled on the routing device, you can use this statement to configure an IS-IS metric that is used exclusively for traffic engineering.</p>
Options	<p>metric—Metric value.</p> <p>Range: 1 through 16,777,215</p> <p>Default: Value of the IGP metric</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 Link and Node Protection for IS-IS Routes on page 141 • metric on page 262 • wide-metrics-only on page 299

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 on page 131• Example: Configuring IS-IS Multicast Topology on page 109

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>

IS-IS Protocol-Specific Tracing Flags

- **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	<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"> • Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces on page 194 • Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces on page 189 • Example: Enabling Packet Checksums on IS-IS Interfaces on page 50

traffic-engineering (Protocols IS-IS)

Syntax	<pre> traffic-engineering { disable; credibility-protocol-preference; family inet { shortcuts { multicast-rpf-routes; } } family inet6 { shortcuts; } multipath { lsp-equal-cost; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the family statement introduced in Junos OS Release 9.3.</p> <p>Support for the credibility-protocol-preference statement introduced in Junos OS Release 9.4.</p> <p>Support for the multipath statement introduced in Junos OS Release 9.6.</p> <p>Support for the lsp-equal-cost statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure traffic engineering properties for IS-IS.</p> <p>IS-IS always performs shortest-path-first (SPF) calculations to determine next hops. For prefixes reachable through a particular next hop, IS-IS places that next hop for that prefix in the inet.0 routing table. In addition, for routers running MPLS, IS-IS installs the prefix for IPv4 routes in the inet.3 routing table as well. The inet.3 table, which is present on the ingress router, contains the host address of each MPLS label-switched path (LSP) egress router. BGP uses this routing table to resolve next-hop addresses.</p> <p>If you enable IS-IS traffic engineering shortcuts and if there is a label-switched path to a point along the path to that prefix, IS-IS installs the prefix in the inet.3 routing table and uses the LSP as a next hop. The net result is that for BGP egress routers for which there is no LSP, BGP automatically uses an LSP along the path to reach the egress router.</p> <p>In Junos OS Release 9.3 and later, IS-IS traffic engineering shortcuts support IPv6 routes. LSPs to be used for shortcuts continue to be signaled using IPv4. However, by default, shortcut routes calculated through IPv6 routes are added to the inet6.3 routing table. The default behavior is for only BGP to use LSPs in its calculations. If you configure MPLS so that both BGP and interior gateway protocols use LSPs for forwarding traffic, shortcut routes calculated through IPv6 are added to the inet6.0 routing table. IS-IS ensures that the IPv6 routes running over the IPv4 MPLS LSP are correctly de-encapsulated at the</p>

tunnel egress by pushing an extra IPv6 explicit null label between the IPv6 payload and the IPv4 transport label.

RSVP LSPs with a higher preference than IS-IS routes are not considered during the computation of traffic engineering shortcuts.

To configure IS-IS so that it uses LSPs as shortcuts when installing information in the inet.3 or inet6.3 routing table, include the following statements:

```
family inet {
  shortcuts {
    multicast-rpf-routes;
  }
}
family inet6 {
  shortcuts;
}
```

For IPv4 traffic, include the **inet** statement. For IPv6 traffic, include the **inet6** statement.

To configure load balancing across multiple LSPs, include the **multipath** statement.

When traffic engineering shortcuts are used, RSVP first looks at the **metric2** value, which is derived from the IGP cost. After this, RSVP considers the LSP metric value. So, if a certain path changes for an LSP and the cost changes, not all LSPs are used to load-balance the network.

When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) combinations can be lost during route calculation. To override this default behavior for load balancing, include the **lsp-equal-cost** statement to retain the equal cost path information in the routing table.

```
multipath {
  lsp-equal-cost;
}
```

Because the inet.3 routing table is present only on ingress routers, you can configure LSP shortcuts only on these routers.

Default IS-IS traffic engineering support is enabled.

By default, IS-IS supports traffic engineering by exchanging basic information with the traffic engineering database. To disable this support, and to disable IS-IS shortcuts if they are configured, include the **disable** statement.

Options **credibility-protocol-preference**—Specify that IS-IS should use the configured protocol preference for IGP routes to determine the traffic engineering database credibility value. By default, the traffic engineering database prefers IS-IS routes even when the routes of another IGP are configured with a lower, that is, more preferred value. Use this statement to override this default behavior.

The traffic engineering database assigns a credibility value to each IGP and prefers the routes of the IGP with the highest credibility value. In Junos OS Release 9.4 and later, you can configure IS-IS to take protocol preference into account to determine the traffic engineering database credibility value. When protocol preference is used to determine the credibility value, IS-IS routes are not automatically preferred by the traffic engineering database, depending on your configuration. 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, 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.



NOTE: This feature is also supported for OSPFv2.

lsp-equal-cost—Configure LSPs to be retained as equal cost paths for load balancing when a better path metric is found during the IS-IS internal routing table calculation. When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) combinations can be lost during route calculation. To override this default IS-IS behavior, include the **lsp-equal-cost** statement for load balancing, so that the equal cost path information is retained in the routing table.

multipath—Enable load balancing for multiple LSPs.

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: Enabling OSPF Traffic Engineering Support*
- [Example: Enabling IS-IS Traffic Engineering Support on page 158](#)
- *traffic-engineering (OSPF)*

wide-metrics-only

Syntax	wide-metrics-only;
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 IS-IS to generate metric values greater than 63 on a per IS-IS level basis. 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 ($2^{24} - 1$). 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.
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 Engineering on page 180 • te-metric on page 291

PART 3

Administration

- [Operational Commands on page 303](#)

CHAPTER 13

Operational Commands

- clear bfd adaptation
- clear bfd session
- clear isis adjacency
- clear isis database
- clear isis overload
- clear isis statistics
- ping clns
- restart
- show bfd session
- show isis adjacency
- show isis authentication
- show isis backup coverage
- show isis backup label-switched-path
- show isis backup spf results
- show isis context-identifier
- show isis database
- show isis hostname
- show isis interface
- show isis overview
- show isis route
- show isis spf
- show isis statistics
- show policy
- show policy conditions
- show route
- show route active-path
- show route advertising-protocol
- show route all

- `show route best`
- `show route brief`
- `show route detail`
- `show route exact`
- `show route export`
- `show route extensive`
- `show route forwarding-table`
- `show route hidden`
- `show route inactive-path`
- `show route instance`
- `show route next-hop`
- `show route output`
- `show route protocol`
- `show route receive-protocol`
- `show route table`
- `show route terse`
- `show security keychain`
- `test policy`
- `traceroute clns`

clear bfd adaptation

Syntax	clear bfd adaptation <address <i>session-address</i> > <discriminator <i>discr-number</i> >
Release Information	Command introduced before Junos OS Release 7.4.
Description	<p>Clear adaptation for Bidirectional Forwarding Detection (BFD) sessions. BFD is a simple hello mechanism that detects failures in a network. Configured BFD interval timers can change, adapting to network situations. Use this command to return BFD interval timers to their configured values.</p> <p>The clear bfd adaptation command is hitless, meaning that the command does not affect traffic flow on the routing device.</p>
Options	<p>none—Clear adaptation for all BFD sessions.</p> <p>address <i>session-address</i>—(Optional) Clear adaptation for all BFD sessions matching the specified address.</p> <p>discriminator <i>discr-number</i>—(Optional) Clear adaptation for the local BFD session matching the specified discriminator.</p>
Additional Information	For more information, see the description of the bfd-liveness-detection configuration statement in the <i>Junos Routing Protocols Configuration Guide</i> .
Required Privilege Level	clear
List of Sample Output	clear bfd adaptation on page 305
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bfd adaptation

```
user@host> clear bfd adaptation
```

clear bfd session

Syntax	<code>clear bfd session</code> <code><address <i>session-address</i>></code> <code><discriminator <i>discr-number</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switch and QFX Series)	<code>clear bfd session</code> <code><address <i>session-address</i>></code> <code><discriminator <i>discr-number</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Drop one or more Bidirectional Forwarding Detection (BFD) sessions.
Options	none —Drop all BFD sessions. address <i>session-address</i> —(Optional) Drop all BFD sessions matching the specified address. discriminator <i>discr-number</i> —(Optional) Drop the local BFD session matching the specified discriminator. 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 bfd session on page 327
List of Sample Output	clear bfd session on page 306
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bfd session

```
user@host> clear bfd session
```


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	<p>none—Remove all entries from the adjacency database.</p> <p>instance <i>instance-name</i>—(Optional) Clear all adjacencies for the specified routing instance only.</p> <p>interface <i>interface-name</i>—(Optional) Clear all adjacencies for 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> <p>neighbor—(Optional) Clear adjacencies for the specified neighbor only.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show isis adjacency on page 335
List of Sample Output	clear isis adjacency on page 307
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>)>
Syntax (EX Series Switches and QFX Series)	clear isis database <entries> <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	Remove the entries from the IS-IS link-state database, which contains prefixes and topology information.
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.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show isis database on page 351
List of Sample Output	clear isis database on page 309
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	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. 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.
Options	none —Reset the IS-IS dynamic overload bit. instance <i>instance-name</i> —(Optional) Reset the IS-IS dynamic overload bit 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	clear
Related Documentation	<ul style="list-style-type: none"> • show isis database on page 351
List of Sample Output	clear isis overload on page 311
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	<p>none—Set IS-IS traffic statistics to zero for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show isis statistics on page 381
List of Sample Output	clear isis statistics on page 313
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

ping clns

Syntax	<pre>ping clns <i>host</i> <brief> <count <i>requests</i>> <detail> <do-not-fragment> <interval <i>seconds</i>> <logical-system <i>logical-system-name</i>> <no-resolve> <pattern <i>string</i>> <rapid> <routing-instance <i>routing-instance-name</i>> <size <i>bytes</i>> <source <i>source-address</i>> <ttl <i>value</i>> <verbose> <wait <i>seconds</i>></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Check the reachability of a remote Connectionless Network Service (CLNS) node. Type Ctrl+c to interrupt a ping clns command.
Options	<p>host—IP address or hostname of the remote system to ping.</p> <p>brief—(Optional) Display brief information.</p> <p>count <i>requests</i>—(Optional) Number of ping requests to send. The range of values is 1 through 2,000,000,000. The default is an unlimited number of requests.</p> <p>detail—(Optional) Include in the output the interface on which the ping reply was received.</p> <p>do-not-fragment—(Optional) Set the do-not-fragment (DF) bit in the IP header of the ping packets.</p> <p>interval <i>seconds</i>—(Optional) How often to send ping requests. The range of values, in seconds, is 1 through infinity. The default value is 1.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Run the command from the specified logical system.</p> <p>no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.</p> <p>pattern <i>string</i>—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.</p> <p>rapid—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of request, include the count option.</p>

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the ping attempt.

size *bytes*—(Optional) Size of ping request packets. The range of values, in bytes, is 0 through 65,468. The default value is 56, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface.

ttl *value*—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is 0 through 255.

verbose—(Optional) Display detailed output.

wait *seconds*—(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is 10 seconds. If this option is used without the count option, a default count of 5 packets is used.

Required Privilege Level network

List of Sample Output [ping clns on page 316](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output


ping clns

```
user@host> ping clns 47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00
PING 47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00
(47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00): 55 data bytes
64 bytes from 47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00: seq=0 ttl=30
time=15.051 ms
64 bytes from 47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00: seq=1 ttl=30
time=10.370 ms
64 bytes from 47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00: seq=2 ttl=30
time=10.367 ms
--- ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 10.367/11.929/15.051/2.207 ms
```

restart

Syntax	<pre>restart <adaptive-services ancpd-service application-identification audit-process auto-configuration captive-portal-content-delivery ce-l2tp-service chassis-control class-of-service clksyncd-service database-replication datapath-trace-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging ethernet-connectivity-fault-management ethernet-link-fault-management event-processing firewall general-authentication-service gracefully iccp-service idp-policy immediately interface-control ipsec-key-management kernel-replication l2-learning l2cpd-service l2tp-service l2tp-universal-edge lacp license-service link-management local-policy-decision-function mac-validation mib-process mobile-ip mountd-service mpls-traceroute mspd multicast-snooping named-service nfsd-service packet-triggered-subscribers peer-selection-service pgcp-service pgm pic-services-logging pki-service ppp ppp-service pppoe protected-system-domain-service redundancy-interface-process remote-operations root-system-domain-service routing <logical-system <i>logical-system-name</i>> sampling sbc-configuration-process sdk-service service-deployment services services pgcp gateway <i>gateway-name</i> snmp soft static-subscribers statistics-service subscriber-management subscriber-management-helper tunnel-oamd usb-control vrrp web-management> <gracefully immediately soft></pre>
Syntax (ACX Series Routers)	<pre>restart <adaptive-services audit-process auto-configuration autoinstallation chassis-control class-of-service clksyncd-service database-replication dhcp-service diameter-service disk-monitoring dynamic-flow-capture ethernet-connectivity-fault-management ethernet-link-fault-management event-processing firewall general-authentication-service gracefully immediately interface-control ipsec-key-management l2-learning lacp link-management mib-process mobile-ip mountd-service mpls-traceroute mspd named-service nfsd-service pgm pki-service ppp pppoe redundancy-interface-process remote-operations routing sampling sdk-service secure-neighbor-discovery service-deployment services snmp soft statistics-service subscriber-management subscriber-management-helper tunnel-oamd vrrp></pre>
Syntax (EX Series Switches)	<pre>restart <autoinstallation chassis-control class-of-service database-replication dhcp dhcp-service diameter-service dot1x-protocol ethernet-link-fault-management ethernet-switching event-processing firewall general-authentication-service interface-control kernel-replication l2-learning lacp license-service link-management lldpd-service mib-process mountd-service multicast-snooping pgm redundancy-interface-process remote-operations routing secure-neighbor-discovery service-deployment sflow-service snmp vrrp web-management></pre>
Syntax (Routing Matrix)	<pre>restart <adaptive-services audit-process chassis-control class-of-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i>> sampling service-deployment snmp> <all all-lcc lcc <i>number</i>></pre>

	<gracefully immediately soft>
Syntax (J Series Routing Platform)	<p>restart</p> <p><adaptive-services audit-process chassis-control class-of-service dhcp dialer-services dlsw event-processing firewall interface-control ipsec-key-management isdn-signaling l2-learning l2tp-service mib-process network-access-service pgm ppp pppoe remote-operations routing <logical-system <i>logical-system-name</i>> sampling service-deployment snmp usb-control web-management></p> <p><gracefully immediately soft></p>
Syntax (TX Matrix Routers)	<p>restart</p> <p><adaptive-services audit-process chassis-control class-of-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i>> sampling service-deployment snmp statistics-service></p> <p><all-chassis all-lcc lcc <i>number</i> scc></p> <p><gracefully immediately soft></p>
Syntax (TX Matrix Plus Routers)	<p>restart</p> <p><adaptive-services audit-process chassis-control class-of-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging event-processing firewall interface-control ipsec-key-management kernel-replication l2-learning l2tp-service lacp link-management mib-process pgm pic-services-logging ppp pppoe redundancy-interface-process remote-operations routing <logical-system <i>logical-system-name</i>> sampling service-deployment snmp statistics-service></p> <p><all-chassis all-lcc all-sfc lcc <i>number</i> sfc <i>number</i>></p> <p><gracefully immediately soft></p>
Syntax (MX Series Routers)	<p>restart</p> <p><adaptive-services ancpd-service application-identification audit-process auto-configuration captive-portal-content-delivery ce-l2tp-service chassis-control class-of-service clksyncd-service database-replication datapath-trace-service dhcp-service diameter-service disk-monitoring dynamic-flow-capture ecc-error-logging ethernet-connectivity-fault-management ethernet-link-fault-management event-processing firewall general-authentication-service gracefully iccp-service idp-policy immediately interface-control ipsec-key-management kernel-replication l2-learning l2cpd-service l2tp-service l2tp-universal-edge lacp license-service link-management local-policy-decision-function mac-validation mib-process mobile-ip mounstd-service mpls-traceroute mspd multicast-snooping named-service nfsd-service packet-triggered-subscribers peer-selection-service pgcp-service pgm pic-services-logging pki-service ppp ppp-service pppoe protected-system-domain-service redundancy-interface-process remote-operations root-system-domain-service routing routing <logical-system <i>logical-system-name</i>> sampling sbc-configuration-process sdk-service service-deployment services services pgcp gateway <i>gateway-name</i> snmp soft static-subscribers statistics-service subscriber-management subscriber-management-helper tunnel-oamd usb-control vrrp web-management></p> <p><all-members></p> <p><gracefully immediately soft></p> <p><local></p> <p><member <i>member-id</i>></p>

Syntax (J Series Routers)	<pre>restart <adaptive-services audit-process chassis-control class-of-service dhcp dhcp-service dialer-services diameter-service dlsf event-processing firewall interface-control ipsec-key-management isdn-signaling l2ald l2-learning l2tp-service mib-process network-access-service pgm ppp pppoe remote-operations routing <logical-system logical-system-name> sampling service-deployment snmp usb-control web-management> <gracefully immediately soft></pre>
Syntax (QFX Series)	<pre>restart <adaptive-services audit-process chassis-control class-of-service dialer-services diameter-service dlsf ethernet-connectivity event-processing fibre-channel firewall general-authentication-service igmp-host-services interface-control ipsec-key-management isdn-signaling l2ald l2-learning l2tp-service mib-process named-service network-access-service nstrace-process pgm ppp pppoe redundancy-interface-process remote-operations logical-system-name> routing sampling secure-neighbor-discovery service-deployment snmp usb-control web-management> <gracefully immediately soft></pre>
Release Information	<p>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 11.1 for the QFX Series. Command introduced in Junos OS Release 12.2 for ACX Series routers. Options added:</p> <ul style="list-style-type: none"> • dynamic-flow-capture in Junos OS Release 7.4. • dlsf in Junos OS Release 7.5. • event-processing in Junos OS Release 7.5. • ppp in Junos OS Release 7.5. • l2ald in Junos OS Release 8.0. • link-management in Release 8.0. • pgcp-service in Junos OS Release 8.4. • sbc-configuration-process in Junos OS Release 9.5. • services pgcp gateway in Junos OS Release 9.6. • sfc and all-sfc for the TX Matrix Router in Junos OS Release 9.6.
Description	Restart a Junos OS process.
	
<p>CAUTION: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.</p>	
Options	none—Same as gracefully .

- adaptive-services**—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.
- all-chassis**—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.
- all-lcc**—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.
- all-members**—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.
- all-sfc**—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).
- ancpd-service**—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.
- application-identification**—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.
- audit-process**—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.
- auto-configuration**—(Optional) Restart the Interface Auto-Configuration process.
- autoinstallation**—(EX Series switches only) (Optional) Restart the autoinstallation process.
- captive-portal-content-delivery**—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.
- ce-l2tp-service**—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.
- chassis-control**—(Optional) Restart the chassis management process.
- class-of-service**—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.
- clksyncd-service**—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

database-replication—(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.

datapath-trace-service—(Optional) Restart the packet path tracing process.

dhcp—(J Series routers and EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

dhcp-service—(Optional) Restart the Dynamic Host Configuration Protocol process.

dialer-services—(J Series routers and EX Series switches only) (Optional) Restart the ISDN dial-out process.

diameter-service—(Optional) Restart the diameter process.

disk-monitoring—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

dls—(J Series routers and QFX Series only) (Optional) Restart the data link switching (DLSw) service.

dot1x-protocol—(EX Series switches only) (Optional) Restart the port-based network access control process.

dynamic-flow-capture—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

ecc-error-logging—(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

ethernet-connectivity-fault-management—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

ethernet-link-fault-management—(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

ethernet-switching—(EX Series switches only) (Optional) Restart the Ethernet switching process.

event-processing—(Optional) Restart the event process (eventd).

fibre-channel—(QFX Series only) (Optional) Restart the Fibre Channel process.

firewall—(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

general-authentication-service—(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.

gracefully—(Optional) Restart the software process.

iccp-service—(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.

idp-policy—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.

immediately—(Optional) Immediately restart the software process.

interface-control—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

ipsec-key-management—(Optional) Restart the IPsec key management process.

isdn-signaling—(J Series routers and QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.

kernel-replication—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

l2-learning—(Optional) Restart the Layer 2 address flooding and learning process.

l2cpd-service—(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

l2tp-service—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

l2tp-universal-edge—(MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

lACP—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

license-service—(EX Series switches only) (Optional) Restart the feature license management process.

link-management— (TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

lldpd-service—(EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

local—(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

local-policy-decision-function— (Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

mac-validation— (Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member *member-id*—(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

mib-process—(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.

mobile-ip—(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.

mountd-service—(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.

mpls-traceroute—(Optional) Restart the MPLS Periodic Traceroute process.

mspd—(Optional) Restart the Multiservice process.

multicast-snooping—(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

named-service—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

network-access-service—(J Series routers and QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.

nfsd-service—(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

packet-triggered-subscribers—(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

peer-selection-service—(Optional) Restart the Peer Selection Service process.

pgcp-service—(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

pgm—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

pic-services-logging—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

pki-service—(Optional) Restart the PKI Service process.

ppp—(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

ppp-service—(Optional) Restart the Universal Edge PPP process, which is the encapsulation protocol process for transporting IP traffic across Universal Edge routers.

pppoe—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

protected-system-domain-service—(Optional) Restart the Protected System Domain (PSD) process.

redundancy-interface-process—(Optional) Restart the ASP redundancy process.

remote-operations—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

root-system-domain-service—(Optional) Restart the Root System Domain (RSD) service.

routing—(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.

routing <logical-system *logical-system-name*>—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

sampling—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

sbc-configuration-process—(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).

scc—(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).

sdk-service—(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

secure-neighbor-discovery—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

sfc *number*—(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

service-deployment—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

services—(Optional) Restart a service.

services pgcp gateway *gateway-name*—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

sflow-service—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

snmp—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

soft—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

static-subscribers—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

statistics-service—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

subscriber-management—(Optional) Restart the Subscriber Management process.

subscriber-management-helper—(Optional) Restart the Subscriber Management Helper process.

tunnel-oamd—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 PDUs across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

usb-control—(J Series routers and MX Series routers only) (Optional) Restart the USB control process.

vrrp—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

web-management—(J Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

Required Privilege Level

reset

Related Documentation

- *Overview of Junos OS CLI Operational Mode Commands*

List of Sample Output [restart interfaces on page 326](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```

show bfd session

Syntax	<pre>show bfd session <brief detail extensive summary> <address address> <client rsvp-oam (brief detail extensive summary) vpls-oam (brief detail extensive instance instance-name summary)> <discriminator discriminator> <logical-system (all logical-system-name)> <prefix address></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bfd session <brief detail extensive summary> <address address> <client rsvp-oam (brief detail extensive summary) vpls-oam (brief detail extensive instance instance-name summary)> <discriminator discriminator> <prefix address></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Options discriminator and address introduced in Junos OS Release 8.2.</p> <p>Option prefix introduced in Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Option client introduced in Junos OS Release 12.3R3.</p>
Description	Display information about active Bidirectional Forwarding Detection (BFD) sessions.
Options	<p>none—(Same as brief) Display information about active BFD sessions.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>address address—(Optional) Display information about the BFD session for the specified neighbor address.</p> <p>client rsvp-oam (brief detail extensive summary) vpls-oam (brief detail extensive instance instance-name summary)—(Optional) Display information about RSVP-OAM or VPLS-OAM BFD sessions in the specified level of output. For VPLS-OAM, display the specified level of output or display information about all of the BFD sessions for the specified VPLS routing instance.</p> <p>discriminator discriminator—(Optional) Display information about the BFD session using the specified local discriminator.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix address—(Optional) Display information about all of the BFD sessions for the specified LDP forwarding equivalence class (FEC).</p>

Required Privilege Level view

- Related Documentation**
- [clear bfd session on page 306](#)
 - *Examples: Configuring BFD for Static Routes*
 - *Example: Configuring BFD for OSPF*
 - *Example: Configuring BFD for BGP*
 - *Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol*
 - [Example: Configuring BFD for IS-IS on page 91](#)

List of Sample Output

[show bfd session on page 331](#)
[show bfd session brief on page 332](#)
[show bfd session detail on page 332](#)
[show bfd session detail \(with Authentication\) on page 332](#)
[show bfd session address extensive on page 332](#)
[show bfd session client rsvp-oam on page 333](#)
[show bfd session client vpls-oam summary on page 333](#)
[show bfd session client vpls-oam instance instance-name on page 333](#)
[show bfd session extensive on page 333](#)
[show bfd session extensive \(with Authentication\) on page 334](#)
[show bfd session summary on page 334](#)

Output Fields [Table 8 on page 328](#) describes the output fields for the **show bfd session** command. Output fields are listed in the approximate order in which they appear.

Table 8: show bfd session Output Fields

Field Name	Field Description	Level of Output
Address	Address on which the BFD session is active.	brief detail extensive none
State	State of the BFD session: Up , Down , Init (initializing), or Failing .	brief detail extensive none
Interface	Interface on which the BFD session is active.	brief detail extensive none
Detect Time	Negotiated time interval, in seconds, used to detect BFD control packets.	brief detail extensive none
Transmit Interval	Time interval, in seconds, used by the transmitting system to send BFD control packets.	brief detail extensive none
Multiplier	Negotiated multiplier by which the time interval is multiplied to determine the detection time for the transmitting system.	detail extensive
Session up time	How long a BFD session has been established.	detail extensive

Table 8: show bfd session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Client	Protocol for which the BFD session is active: ISIS , OSPF , or Static .	detail extensive
TX interval	Time interval, in seconds, used by the host system to transmit BFD control packets.	brief detail extensive none
RX interval	Time interval, in seconds, used by the host system to receive BFD control packets.	brief detail extensive none
Authenticate	Indicates that BFD authentication is configured.	detail extensive
keychain	Name of the security authentication keychain being used by a specific client. BFD authentication information for a client is provided in a single line and includes the keychain , algo , and mode parameters. Multiple clients can be configured on a BFD session.	extensive
algo	BFD authentication algorithm being used for a specific client: keyed-md5 , keyed-sha-1 , meticulous-keyed-md5 , meticulous-keyed-sha-1 , or simple-password . BFD authentication information for a client is provided in a single line and includes the keychain , algo , and mode parameters. Multiple clients can be configured on a BFD session.	extensive
mode	Level of BFD authentication enforcement being used by a specific client: strict or loose . Strict enforcement indicates that authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated. BFD authentication information for a client is provided in a single line and includes the keychain , algo , and mode parameters. Multiple clients can be configured on a BFD session.	extensive
Local diagnostic	Local diagnostic information about failing BFD sessions.	detail extensive
Remote diagnostic	Remote diagnostic information about failing BFD sessions.	detail extensive
Remote state	Reports whether the remote system's BFD packets have been received and whether the remote system is receiving transmitted control packets.	detail extensive
Version	BFD version: 0 or 1 .	extensive
Replicated	The replicated flag appears when nonstop routing or graceful Routing Engine switchover is configured and the BFD session has been replicated to the backup Routing Engine.	detail extensive
Min async interval	Minimum amount of time, in seconds, between asynchronous control packet transmissions across the BFD session.	extensive
Min slow interval	Minimum amount of time, in seconds, between synchronous control packet transmissions across the BFD session.	extensive

Table 8: show bfd session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Adaptive async TX interval	Transmission interval being used because of adaptation.	extensive
RX interval	Minimum required receive interval.	extensive
Local min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the local system.	extensive
Local min RX interval	Minimum amount of time, in seconds, between control packet detections on the local system.	extensive
Remote min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the remote system.	extensive
Remote min RX interval	Minimum amount of time, in seconds, between control packet detections on the remote system.	extensive
Threshold transmission interval	Threshold for notification if the transmission interval increases.	extensive
Threshold for detection time	Threshold for notification if the detection time increases.	extensive
Local discriminator	Authentication code used by the local system to identify that BFD session.	extensive
Remote discriminator	Authentication code used by the remote system to identify that BFD session.	extensive
Echo mode	Information about the state of echo transmissions on the BFD session.	extensive
Prefix	LDP FEC address associated with the BFD session.	All levels
Egress, Destination	Displays the LDP FEC destination address. This field is displayed only on a router at the egress of an LDP FEC, where the BFD session has an LDP Operation, Administration, and Maintenance (OAM) client.	All levels
Remote is control-plane independent	<p>The BFD session on the remote peer is running on its Packet Forwarding Engine. In this case, when the remote node undergoes a graceful restart, the local peer can help the remote peer with the graceful restart.</p> <p>The following BFD sessions are not distributed to the Packet Forwarding Engine: tunnel-encapsulated sessions, and sessions over integrated routing and bridging (IRB) interfaces.</p>	extensive

Table 8: show bfd session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Authentication	<p>Summary status of BFD authentication:</p> <ul style="list-style-type: none"> • status—enabled/active indicates authentication is configured and active. enabled/inactive indicates authentication is configured but not active. This only occurs when the remote end of the session does not support authentication and loose checking is configured. • keychain—Name of the security authentication keychain associated with the specified BFD session. • algo—BFD authentication algorithm being used: keyed-md5, keyed-sha-1, meticulous-keyed-md5, meticulous-keyed-sha-1, or simple-password. • mode—Level of BFD authentication enforcement: strict or loose. Strict enforcement indicates authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated. <p>This information is only shown if BFD authentication is configured.</p>	extensive
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).	detail extensive
sessions	Total number of active BFD sessions.	All levels
clients	Total number of clients that are hosting active BFD sessions.	All levels
Cumulative transmit rate	Total number of BFD control packets transmitted per second on all active sessions.	All levels
Cumulative receive rate	Total number of BFD control packets received per second on all active sessions.	All levels
Multi-hop, min-recv-TTL	Minimum time to live (TTL) accepted if the session is configured for multihop.	extensive
route table	Route table used if the session is configured for multihop.	extensive
local address	<p>Local address of the source used if the session is configured for multihop.</p> <p>The source IP address for outgoing BFD packets from the egress side of an MPLS BFD session is based on the outgoing interface IP address.</p>	extensive

Sample Output

show bfd session

```

user@host> show bfd session

Address          State   Interface   Detect Time   Transmit Interval   Multiplier
10.9.1.33        Up     so-7/1/0.0   0.600        0.200          3
10.9.1.29        Up     ge-4/0/0.0   0.600        0.200          3

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

```

show bfd session brief

The output for the **show bfd session brief** command is identical to that for the **show bfd session** command. For sample output, see [show bfd session on page 331](#).

show bfd session detail

```
user@host> show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.9.1.33	Up	so-7/1/0.0	0.600	0.200	3
Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3					
Session up time 3d 00:34					
Local diagnostic None, remote diagnostic None					
Remote state Up, version 1					
Replicated					
10.9.1.29	Up	ge-4/0/0.0	0.600	0.200	3
Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3					
Session up time 3d 00:29, previous down time 00:00:01					
Local diagnostic NbrSignal, remote diagnostic AdminDown					
Remote state Up, version 1					

2 sessions, 2 clients

Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

show bfd session detail (with Authentication)

```
user@host> show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.9.1.33	Up	so-7/1/0.0	0.600	0.200	3
Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, Authenticate					
Session up time 3d 00:34					
Local diagnostic None, remote diagnostic None					
Remote state Up, version 1					
Replicated					
10.9.1.29	Up	ge-4/0/0.0	0.600	0.200	3
Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3					
Session up time 3d 00:29, previous down time 00:00:01					
Local diagnostic NbrSignal, remote diagnostic AdminDown					
Remote state Up, version 1					

2 sessions, 2 clients

Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

show bfd session address extensive

```
user@host> show bfd session 10.255.245.212 extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.255.245.212	Up		1.200	0.400	3
Client Static, TX interval 0.400, RX interval 0.400, multiplier 3					
Session up time 00:17:03, previous down time 00:00:14					
Local diagnostic CtlExpire, remote diagnostic NbrSignal					
Remote state Up, version 1					
Replicated					
Min async interval 0.400, min slow interval 1.000					
Adaptive async tx interval 0.400, rx interval 0.400					
Local min tx interval 0.400, min rx interval 0.400, multiplier 3					
Remote min tx interval 0.400, min rx interval 0.400, multiplier 3					

Threshold transmission interval 0.000, Threshold for detection time 0.000
 Local discriminator 6, remote discriminator 16
 Echo mode disabled/inactive
 Multi-hop, min-recv-TTL 255, route-table 0, local-address 10.255.245.205

1 sessions, 1 clients
 Cumulative transmit rate 2.5 pps, cumulative receive rate 2.5 pps

show bfd session client rsvp-oam

user@host> show bfd session client rsvp-oam

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.168.0.223	Up		540.000	180.000	3

1 Up sessions, 0 Down sessions
 1 sessions, 1 clients
 Cumulative transmit rate 0.0 pps, cumulative receive rate 0.0 pps

show bfd session client vpls-oam summary

user@host> show bfd session client vpls-oam summary

1 Up sessions, 1 Down sessions
 2 sessions, 2 clients
 Cumulative transmit rate 2.0 pps, cumulative receive rate 1.0 pps

show bfd session client vpls-oam instance instance-name

user@host> show bfd session client vpls-oam instance vpls

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
127.0.0.1	Up	ae9.0	3.000	1.000	3

1 Up Sessions, 0 Down Sessions
 1 sessions, 1 clients
 Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

show bfd session extensive

user@host> show bfd session extensive

10.31.1.2 Up ge-2/1/8.0 0.030 0.010 3
 Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010
 Session up time 00:10:13
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated
 Min async interval 0.010, min slow interval 1.000
 Adaptive async TX interval 0.010, RX interval 0.010
 Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3
 Remote min TX interval 0.010, min RX interval 0.010, multiplier 3
 Local discriminator 12, remote discriminator 4
 Echo mode disabled/inactive
 Remote is control-plane independent
 Session ID: 0x201
 Micro-BFD Session

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.31.2.2	Up	ge-2/1/4.0	0.030	0.010	3

Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010

```

Session up time 00:10:14
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 0.010, min slow interval 1.000
Adaptive async TX interval 0.010, RX interval 0.010
Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3
Remote min TX interval 0.010, min RX interval 0.010, multiplier 3
Local discriminator 13, remote discriminator 5
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x202

```

```

2 sessions, 2 clients
Cumulative transmit rate 200.0 pps, cumulative receive rate 200.0 pps

```

show bfd session extensive (with Authentication)

```

user@host> show bfd session extensive

```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.168.208.26	Up	so-1/0/0.0	2.400	0.800	10

```

Client Static, TX interval 0.600, RX interval 0.600, Authenticate
    keychain bfd, algo keyed-md5, mode loose
Session up time 00:18:07
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 0.600, min slow interval 1.000
Adaptive async TX interval 0.600, RX interval 0.600
Local min TX interval 0.600, minimum RX interval 0.600, multiplier 10
Remote min TX interval 0.800, min RX interval 0.800, multiplier 3
Local discriminator 2, remote discriminator 3
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd, algo keyed-md5, mode loose

1 sessions, 1 clients
Cumulative transmit rate 1.2 pps, cumulative receive rate 1.2 pps

```

show bfd session summary

```

user@host> show bfd session summary
2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

```

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 307 	
List of Sample Output	<p>show isis adjacency on page 337</p> <p>show isis adjacency brief on page 337</p> <p>show isis adjacency detail on page 338</p> <p>show isis adjacency extensive on page 338</p>	
Output Fields	<p>Table 9 on page 335 describes the output fields for the show isis adjacency command. Output fields are listed in the approximate order in which they appear.</p>	

Table 9: show isis adjacency Output Fields

Field Name	Field Description	Level of Output
Interface	Interface through which the neighbor is reachable.	All levels

Table 9: 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 routing device 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 routing device's outbound link-state PDUs.	detail extensive
IP addresses	IP address of this neighbor.	detail extensive

Table 9: 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 337](#).

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        Seenself
```


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 340 show isis authentication (With Hitless Authentication Key Rollover Configured) on page 340
Output Fields	Table 10 on page 339 describes the output fields for the show isis authentication command. Output fields are listed in the approximate order in which they appear.

Table 10: 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 10: 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	show isis backup coverage <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis backup coverage <instance <i>instance-name</i> >
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 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 on page 141 • show isis backup label-switched-path on page 343
List of Sample Output	show isis backup coverage on page 342
Output Fields	Table 11 on page 341 lists the output fields for the show isis backup coverage command. Output fields are listed in the approximate order in which they appear.

Table 11: 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 11: 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 on page 141 • show isis backup coverage on page 341
List of Sample Output	show isis backup label-switched-path on page 344
Output Fields	Table 12 on page 343 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 12: 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 routing device can detect RSVP hello messages from the neighbor. • Down—The routing device 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 12: 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 on page 141 • show isis backup coverage on page 341
List of Sample Output	<p>show isis backup spf results on page 346</p> <p>show isis backup spf results no-coverage on page 347</p>
Output Fields	<p>Table 13 on page 346 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 13: 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:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0

```



```

        track-item: pro-bng3-j.00-00
        track-item: pro-bng3-i.00-00
        Not eligible, IPV4, Reason: Path loops
        Not eligible, IPV6, Reason: Path loops
    2 nodes

IS-IS level 2 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
2 nodes

```

show isis backup spf results no-coverage

```

user@host> show isis backup spf results no-coverage
IS-IS level 1 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
2 nodes

IS-IS level 2 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21

```

```
Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
2 nodes
```

show isis context-identifier

Syntax	show isis context-identifier <brief detail extensive> <identifier name> <instance instance-name> <logical-system (all logical-system-name)>
Release Information	Command introduced in Junos OS Release 10.4.
Description	Display IS-IS context identifier information.
Options	brief detail extensive —(Optional) Display the specified level of output. identifier name —(Optional) Display information about the specified context identifier. instance instance-name —(Optional) Display entries for the specified routing instance. logical-system (all logical-system-name) —(Optional) Display the context identifier information for all logical systems or for a particular logical system.
Required Privilege Level	View
Output Fields	Table 14 on page 349 lists the output fields for the show isis context-identifier command. Output fields are listed in the approximate order in which they appear.

Table 14: show isis context-identifier Output Fields

Field Name	Field Description	Level of Output
Context	IPv4 address that defines a protection pair. The context is manually configured on both primary and protector PEs.	detail
Owner	Protocol that requires the context.	detail
Role	Role of the PE, which is either primary or protector.	detail
Primary	Name of the primary PE.	detail
Metric	Advertised interior gateway protocol (IGP) metric.	detail

Sample Output

```
user@host> show isis context-identifier detail
```

```
IS-IS context database:
```

```
Context      Owner      Role      Primary      Metric
2.2.4.3      MPLS      Primary   pro3-e        1
  Advertiser pro3-e, Router ID 10.255.245.198, Metric 1, Level 1
  Advertiser pro3-e, Router ID 10.255.245.198, Metric 1, Level 2
  Advertiser pro3-c, Router ID 10.255.245.196, Metric 11, Level 2
```


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 309
List of Sample Output	<p>show isis database on page 353</p> <p>show isis database brief on page 354</p> <p>show isis database detail on page 354</p> <p>show isis database extensive on page 354</p>

Output Fields Table 15 on page 352 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 15: 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 15: 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, IP address, local interface index, and remote interface index. <p>The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.</p>	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 353](#).

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 extensive
```

IS-IS level 1 link-state database:

```
Router-A.00-00 Sequence: 0x1, Checksum: 0xf75c, Lifetime: 1116 secs
```


IP prefix: 192.168.0.1/32 Metric: 0 Internal Up

Header: LSP ID: Router-A.00-00, Length: 85 bytes
 Allocated length: 1492 bytes, Router ID: 192.168.0.1
 Remaining lifetime: 1116 secs, Level: 1, Interface: 0
 Estimated free bytes: 1353, Actual free bytes: 1407
 Aging timer expires in: 1116 secs
 Protocols: IP, IPv6

Packet: LSP ID: Router-A.00-00, Length: 85 bytes, Lifetime : 1200 secs
 Checksum: 0xf75c, Sequence: 0x1, Attributes: 0x3 <L1 L2>
 NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
 Packet type: 18, Packet version: 1, Max area: 0

TLVs:

Area address: 49.0002 (3)
 LSP Buffer Size: 1492
 Speaks: IP
 Speaks: IPV6
 IP router id: 192.168.0.1
 IP address: 192.168.0.1
 Hostname: Router-A
 IP prefix: 192.168.0.1/32, Internal, Metric: default 0, Up
 IP extended prefix: 192.168.0.1/32 metric 0 up
 No queued transmissions

IS-IS level 2 link-state database:

Router-A.00-00 Sequence: 0x5, Checksum: 0x3196, Lifetime: 1158 secs
 IS neighbor: Router-B.02 Metric: 10
 Two-way fragment: Router-B.02-00, Two-way first fragment: Router-B.02-00
 IS neighbor: Router-E.02 Metric: 10
 Two-way fragment: Router-E.02-00, Two-way first fragment: Router-E.02-00
 IP prefix: 10.0.0.0/30 Metric: 10 Internal Up
 IP prefix: 10.0.0.4/30 Metric: 10 Internal Up
 IP prefix: 192.168.0.1/32 Metric: 0 Internal Up

Header: LSP ID: Router-A.00-00, Length: 208 bytes
 Allocated length: 1492 bytes, Router ID: 192.168.0.1
 Remaining lifetime: 1158 secs, Level: 2, Interface: 0
 Estimated free bytes: 1233, Actual free bytes: 1284
 Aging timer expires in: 1158 secs
 Protocols: IP, IPv6

Packet: LSP ID: Router-A.00-00, Length: 208 bytes, Lifetime : 1198 secs
 Checksum: 0x3196, Sequence: 0x5, 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: 49.0002 (3)
 LSP Buffer Size: 1492
 Speaks: IP
 Speaks: IPV6
 IP router id: 192.168.0.1
 IP address: 192.168.0.1
 Hostname: Router-A
 IP prefix: 192.168.0.1/32, Internal, Metric: default 0, Up
 IP prefix: 10.0.0.4/30, Internal, Metric: default 10, Up
 IP prefix: 10.0.0.0/30, Internal, Metric: default 10, Up
 IP extended prefix: 192.168.0.1/32 metric 0 up

```

IP extended prefix: 10.0.0.4/30 metric 10 up
IP extended prefix: 10.0.0.0/30 metric 10 up
IS neighbor: Router-E.02, Internal, Metric: default 10
IS neighbor: Router-B.02, Internal, Metric: default 10
IS extended neighbor: Router-E.02, Metric: default 10
  IP address: 10.0.0.1
    Local interface index: 101, Remote interface index: 0
IS extended neighbor: Router-B.02, Metric: default 10
  IP address: 10.0.0.5
    Local interface index: 102, Remote interface index: 0
No queued transmissions

Router-B.00-00 Sequence: 0x5, Checksum: 0xf8f, Lifetime: 1183 secs
  IS neighbor: Router-B.02                      Metric: 10
    Two-way fragment: Router-B.02-00, Two-way first fragment: Router-B.02-00
  IS neighbor: Router-C.02                      Metric: 10
    Two-way fragment: Router-C.02-00, Two-way first fragment: Router-C.02-00
IP prefix: 10.0.0.4/30                          Metric: 10 Internal Up
IP prefix: 10.0.0.8/30                          Metric: 10 Internal Up
IP prefix: 192.168.0.2/32                      Metric: 0 Internal Up

Header: LSP ID: Router-B.00-00, Length: 208 bytes
  Allocated length: 284 bytes, Router ID: 192.168.0.2
  Remaining lifetime: 1183 secs, Level: 2, Interface: 102
  Estimated free bytes: 114, Actual free bytes: 76
  Aging timer expires in: 1183 secs
  Protocols: IP, IPv6

Packet: LSP ID: Router-B.00-00, Length: 208 bytes, Lifetime : 1196 secs
  Checksum: 0xf8f, Sequence: 0x5, 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: 49.0002 (3)
  LSP Buffer Size: 1492
  Speaks: IP
  Speaks: IPV6
  IP router id: 192.168.0.2
  IP address: 192.168.0.2
  Hostname: Router-B
  IP prefix: 192.168.0.2/32, Internal, Metric: default 0, Up
  IP prefix: 10.0.0.4/30, Internal, Metric: default 10, Up
  IP prefix: 10.0.0.8/30, Internal, Metric: default 10, Up
  IP extended prefix: 192.168.0.2/32 metric 0 up
  IP extended prefix: 10.0.0.4/30 metric 10 up
  IP extended prefix: 10.0.0.8/30 metric 10 up
  IS neighbor: Router-B.02, Internal, Metric: default 10
  IS neighbor: Router-C.02, Internal, Metric: default 10
  IS extended neighbor: Router-B.02, Metric: default 10
    IP address: 10.0.0.6
      Local interface index: 108, Remote interface index: 0
  IS extended neighbor: Router-C.02, Metric: default 10
    IP address: 10.0.0.9
      Local interface index: 109, Remote interface index: 0
No queued transmissions

Router-B.02-00 Sequence: 0x1, Checksum: 0x3c7c, Lifetime: 1156 secs
  IS neighbor: Router-A.00                      Metric: 0
    Two-way fragment: Router-A.00-00, Two-way first fragment: Router-A.00-00
  IS neighbor: Router-B.00                      Metric: 0

```

Two-way fragment: Router-B.00-00, Two-way first fragment: Router-B.00-00

Header: LSP ID: Router-B.02-00, Length: 76 bytes
 Allocated length: 284 bytes, Router ID: 0.0.0.0
 Remaining lifetime: 1156 secs, Level: 2, Interface: 102
 Estimated free bytes: 208, Actual free bytes: 208
 Aging timer expires in: 1156 secs

Packet: LSP ID: Router-B.02-00, Length: 76 bytes, Lifetime : 1196 secs
 Checksum: 0x3c7c, Sequence: 0x1, 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:
 IS neighbor: Router-B.00, Internal, Metric: default 0
 IS neighbor: Router-A.00, Internal, Metric: default 0
 IS extended neighbor: Router-B.00, Metric: default 0
 IS extended neighbor: Router-A.00, Metric: default 0
 No queued transmissions

Router-C.00-00 Sequence: 0x5, Checksum: 0x255b, Lifetime: 1182 secs
 IS neighbor: Router-C.02 Metric: 10
 Two-way fragment: Router-C.02-00, Two-way first fragment: Router-C.02-00
 IS neighbor: Router-D.03 Metric: 10
 Two-way fragment: Router-D.03-00, Two-way first fragment: Router-D.03-00
 IP prefix: 10.0.0.8/30 Metric: 10 Internal Up
 IP prefix: 10.0.0.12/30 Metric: 10 Internal Up
 IP prefix: 192.168.0.3/32 Metric: 0 Internal Up

Header: LSP ID: Router-C.00-00, Length: 208 bytes
 Allocated length: 284 bytes, Router ID: 192.168.0.3
 Remaining lifetime: 1182 secs, Level: 2, Interface: 102
 Estimated free bytes: 114, Actual free bytes: 76
 Aging timer expires in: 1182 secs
 Protocols: IP, IPv6

Packet: LSP ID: Router-C.00-00, Length: 208 bytes, Lifetime : 1196 secs
 Checksum: 0x255b, Sequence: 0x5, 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: 49.0002 (3)
 LSP Buffer Size: 1492
 Speaks: IP
 Speaks: IPV6
 IP router id: 192.168.0.3
 IP address: 192.168.0.3
 Hostname: Router-C
 IP prefix: 192.168.0.3/32, Internal, Metric: default 0, Up
 IP prefix: 10.0.0.8/30, Internal, Metric: default 10, Up
 IP prefix: 10.0.0.12/30, Internal, Metric: default 10, Up
 IP extended prefix: 192.168.0.3/32 metric 0 up
 IP extended prefix: 10.0.0.8/30 metric 10 up
 IP extended prefix: 10.0.0.12/30 metric 10 up
 IS neighbor: Router-C.02, Internal, Metric: default 10
 IS neighbor: Router-D.03, Internal, Metric: default 10
 IS extended neighbor: Router-C.02, Metric: default 10
 IP address: 10.0.0.10
 Local interface index: 105, Remote interface index: 0
 IS extended neighbor: Router-D.03, Metric: default 10

```

    IP address: 10.0.0.13
    Local interface index: 106, Remote interface index: 0
    No queued transmissions

Router-C.02-00 Sequence: 0x1, Checksum: 0xaa09, Lifetime: 1181 secs
  IS neighbor: Router-B.00 Metric: 0
    Two-way fragment: Router-B.00-00, Two-way first fragment: Router-B.00-00
  IS neighbor: Router-C.00 Metric: 0
    Two-way fragment: Router-C.00-00, Two-way first fragment: Router-C.00-00

Header: LSP ID: Router-C.02-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1181 secs, Level: 2, Interface: 102
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1181 secs

Packet: LSP ID: Router-C.02-00, Length: 76 bytes, Lifetime : 1194 secs
  Checksum: 0xaa09, Sequence: 0x1, 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:
  IS neighbor: Router-C.00, Internal, Metric: default 0
  IS neighbor: Router-B.00, Internal, Metric: default 0
  IS extended neighbor: Router-C.00, Metric: default 0
  IS extended neighbor: Router-B.00, Metric: default 0
  No queued transmissions

Router-D.00-00 Sequence: 0x4, Checksum: 0x8ab7, Lifetime: 1180 secs
  IS neighbor: Router-D.02 Metric: 10
    Two-way fragment: Router-D.02-00, Two-way first fragment: Router-D.02-00
  IS neighbor: Router-D.03 Metric: 10
    Two-way fragment: Router-D.03-00, Two-way first fragment: Router-D.03-00
  IP prefix: 10.0.0.12/30 Metric: 10 Internal Up
  IP prefix: 10.0.0.20/30 Metric: 10 Internal Up
  IP prefix: 192.168.0.4/32 Metric: 0 Internal Up

Header: LSP ID: Router-D.00-00, Length: 208 bytes
  Allocated length: 284 bytes, Router ID: 192.168.0.4
  Remaining lifetime: 1180 secs, Level: 2, Interface: 102
  Estimated free bytes: 114, Actual free bytes: 76
  Aging timer expires in: 1180 secs
  Protocols: IP, IPv6

Packet: LSP ID: Router-D.00-00, Length: 208 bytes, Lifetime : 1192 secs
  Checksum: 0x8ab7, Sequence: 0x4, 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: 49.0002 (3)
  LSP Buffer Size: 1492
  Speaks: IP
  Speaks: IPV6
  IP router id: 192.168.0.4
  IP address: 192.168.0.4
  Hostname: Router-D
  IP prefix: 192.168.0.4/32, Internal, Metric: default 0, Up
  IP prefix: 10.0.0.12/30, Internal, Metric: default 10, Up
  IP prefix: 10.0.0.20/30, Internal, Metric: default 10, Up
  IP extended prefix: 192.168.0.4/32 metric 0 up

```

```

IP extended prefix: 10.0.0.12/30 metric 10 up
IP extended prefix: 10.0.0.20/30 metric 10 up
IS neighbor: Router-D.02, Internal, Metric: default 10
IS neighbor: Router-D.03, Internal, Metric: default 10
IS extended neighbor: Router-D.02, Metric: default 10
  IP address: 10.0.0.22
  Local interface index: 115, Remote interface index: 0
IS extended neighbor: Router-D.03, Metric: default 10
  IP address: 10.0.0.14
  Local interface index: 114, Remote interface index: 0
No queued transmissions

Router-D.02-00 Sequence: 0x1, Checksum: 0xebbc, Lifetime: 1128 secs
IS neighbor: Router-D.00 Metric: 0
  Two-way fragment: Router-D.00-00, Two-way first fragment: Router-D.00-00
IS neighbor: Router-F.00 Metric: 0
  Two-way fragment: Router-F.00-00, Two-way first fragment: Router-F.00-00

Header: LSP ID: Router-D.02-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1128 secs, Level: 2, Interface: 101
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1128 secs

Packet: LSP ID: Router-D.02-00, Length: 76 bytes, Lifetime : 1160 secs
  Checksum: 0xebbc, Sequence: 0x1, 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:
  IS neighbor: Router-D.00, Internal, Metric: default 0
  IS neighbor: Router-F.00, Internal, Metric: default 0
  IS extended neighbor: Router-D.00, Metric: default 0
  IS extended neighbor: Router-F.00, Metric: default 0
No queued transmissions

Router-D.03-00 Sequence: 0x1, Checksum: 0x129b, Lifetime: 1180 secs
IS neighbor: Router-C.00 Metric: 0
  Two-way fragment: Router-C.00-00, Two-way first fragment: Router-C.00-00
IS neighbor: Router-D.00 Metric: 0
  Two-way fragment: Router-D.00-00, Two-way first fragment: Router-D.00-00

Header: LSP ID: Router-D.03-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1180 secs, Level: 2, Interface: 101
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1180 secs

Packet: LSP ID: Router-D.03-00, Length: 76 bytes, Lifetime : 1192 secs
  Checksum: 0x129b, Sequence: 0x1, 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:
  IS neighbor: Router-D.00, Internal, Metric: default 0
  IS neighbor: Router-C.00, Internal, Metric: default 0
  IS extended neighbor: Router-D.00, Metric: default 0
  IS extended neighbor: Router-C.00, Metric: default 0
No queued transmissions

Router-E.00-00 Sequence: 0x4, Checksum: 0x9da9, Lifetime: 1155 secs

```

```

IS neighbor: Router-E.02                      Metric:      10
  Two-way fragment: Router-E.02-00, Two-way first fragment: Router-E.02-00
IS neighbor: Router-F.02                      Metric:      20
  Two-way fragment: Router-F.02-00, Two-way first fragment: Router-F.02-00
IP prefix: 10.0.0.0/30                        Metric:      10 Internal Up
IP prefix: 10.0.0.16/30                      Metric:      20 Internal Up
IP prefix: 192.168.0.5/32                    Metric:       0 Internal Up

```

```

Header: LSP ID: Router-E.00-00, Length: 208 bytes
  Allocated length: 284 bytes, Router ID: 192.168.0.5
  Remaining lifetime: 1155 secs, Level: 2, Interface: 101
  Estimated free bytes: 114, Actual free bytes: 76
  Aging timer expires in: 1155 secs
  Protocols: IP, IPv6

```

```

Packet: LSP ID: Router-E.00-00, Length: 208 bytes, Lifetime : 1185 secs
  Checksum: 0x9da9, Sequence: 0x4, 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: 49.0002 (3)
LSP Buffer Size: 1492
Speaks: IP
Speaks: IPV6
IP router id: 192.168.0.5
IP address: 192.168.0.5
Hostname: Router-E
IP prefix: 192.168.0.5/32, Internal, Metric: default 0, Up
IP prefix: 10.0.0.16/30, Internal, Metric: default 20, Up
IP prefix: 10.0.0.0/30, Internal, Metric: default 10, Up
IP extended prefix: 192.168.0.5/32 metric 0 up
IP extended prefix: 10.0.0.16/30 metric 20 up
IP extended prefix: 10.0.0.0/30 metric 10 up
IS neighbor: Router-E.02, Internal, Metric: default 10
IS neighbor: Router-F.02, Internal, Metric: default 20
IS extended neighbor: Router-E.02, Metric: default 10
  IP address: 10.0.0.2
  Local interface index: 112, Remote interface index: 0
IS extended neighbor: Router-F.02, Metric: default 20
  IP address: 10.0.0.17
  Local interface index: 111, Remote interface index: 0
No queued transmissions

```

```

Router-E.02-00 Sequence: 0x1, Checksum: 0xb4fa, Lifetime: 1130 secs
IS neighbor: Router-A.00                      Metric:       0
  Two-way fragment: Router-A.00-00, Two-way first fragment: Router-A.00-00
IS neighbor: Router-E.00                      Metric:       0
  Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00

```

```

Header: LSP ID: Router-E.02-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1130 secs, Level: 2, Interface: 101
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1130 secs

```

```

Packet: LSP ID: Router-E.02-00, Length: 76 bytes, Lifetime : 1161 secs
  Checksum: 0xb4fa, Sequence: 0x1, 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:

IS neighbor: Router-E.00, Internal, Metric: default 0
 IS neighbor: Router-A.00, Internal, Metric: default 0
 IS extended neighbor: Router-E.00, Metric: default 0
 IS extended neighbor: Router-A.00, Metric: default 0

No queued transmissions

Router-F.00-00 Sequence: 0x5, Checksum: 0x94bd, Lifetime: 1153 secs
 IS neighbor: Router-D.02 Metric: 10
 Two-way fragment: Router-D.02-00, Two-way first fragment: Router-D.02-00
 IS neighbor: Router-F.02 Metric: 10
 Two-way fragment: Router-F.02-00, Two-way first fragment: Router-F.02-00
 IP prefix: 10.0.0.16/30 Metric: 10 Internal Up
 IP prefix: 10.0.0.20/30 Metric: 10 Internal Up
 IP prefix: 192.168.0.6/32 Metric: 0 Internal Up

Header: LSP ID: Router-F.00-00, Length: 208 bytes
 Allocated length: 284 bytes, Router ID: 192.168.0.6
 Remaining lifetime: 1153 secs, Level: 2, Interface: 101
 Estimated free bytes: 76, Actual free bytes: 76
 Aging timer expires in: 1153 secs
 Protocols: IP, IPv6

Packet: LSP ID: Router-F.00-00, Length: 208 bytes, Lifetime : 1183 secs
 Checksum: 0x94bd, Sequence: 0x5, 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: 49.0002 (3)
 LSP Buffer Size: 1492
 Speaks: IP
 Speaks: IPV6
 IP router id: 192.168.0.6
 IP address: 192.168.0.6
 Hostname: Router-F
 IP prefix: 192.168.0.6/32, Internal, Metric: default 0, Up
 IP prefix: 10.0.0.16/30, Internal, Metric: default 10, Up
 IP prefix: 10.0.0.20/30, Internal, Metric: default 10, Up
 IP extended prefix: 192.168.0.6/32 metric 0 up
 IP extended prefix: 10.0.0.16/30 metric 10 up
 IP extended prefix: 10.0.0.20/30 metric 10 up
 IS neighbor: Router-D.02, Internal, Metric: default 10
 IS neighbor: Router-F.02, Internal, Metric: default 10
 IS extended neighbor: Router-D.02, Metric: default 10
 IP address: 10.0.0.21
 Local interface index: 94, Remote interface index: 0
 IS extended neighbor: Router-F.02, Metric: default 10
 IP address: 10.0.0.18
 Local interface index: 93, Remote interface index: 0

No queued transmissions

Router-F.02-00 Sequence: 0x1, Checksum: 0xf5ae, Lifetime: 1153 secs
 IS neighbor: Router-E.00 Metric: 0
 Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00
 IS neighbor: Router-F.00 Metric: 0
 Two-way fragment: Router-F.00-00, Two-way first fragment: Router-F.00-00

Header: LSP ID: Router-F.02-00, Length: 76 bytes
 Allocated length: 284 bytes, Router ID: 0.0.0.0
 Remaining lifetime: 1153 secs, Level: 2, Interface: 101

Estimated free bytes: 208, Actual free bytes: 208
Aging timer expires in: 1153 secs

Packet: LSP ID: Router-F.02-00, Length: 76 bytes, Lifetime : 1183 secs
Checksum: 0xf5ae, Sequence: 0x1, 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:

IS neighbor: Router-F.00, Internal, Metric: default 0
IS neighbor: Router-E.00, Internal, Metric: default 0
IS extended neighbor: Router-F.00, Metric: default 0
IS extended neighbor: Router-E.00, Metric: default 0

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. This command displays the system ID-to-name cache. The output shows if the mapping has been learned by receipt of a Hostname TLV #137 (type dynamic) configured in Junos OS with the set system host-name command, or a static mapping defined in Junos OS with the set system static-host-mapping hostname sysid command (type static). The local router always has its type set to static even if static-host-mapping is not configured.
Options	none —Display IS-IS hostname database information. 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 hostname on page 364
Output Fields	Table 16 on page 363 describes the output fields for the show isis hostname command. Output fields are listed in the approximate order in which they appear.

Table 16: 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
1921.6800.4201 isis1
1921.6800.4202 isis2
1921.6800.4203 isis3
```

	Type
1921.6800.4201 isis1	Dynamic
1921.6800.4202 isis2	Static
1921.6800.4203 isis3	Dynamic

show isis interface


Syntax	<pre>show isis interface <brief detail extensive> <interface-name> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis interface <brief detail extensive> <interface-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Display status information about IS-IS-enabled interfaces.</p>
<div>  <p>NOTE: If the configured metric for an IS-IS level is above 63, and the <code>wide-metrics-only</code> statement is not configured, the <code>show isis interface detail</code> command and the <code>show isis interface extensive</code> command display 63 as the metric value for that level. Configure the <code>wide-metrics-only</code> statement to generate metric values greater than 63 on a per IS-IS level basis.</p> <p>The <code>show isis interface</code> 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 logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Enabling Wide IS-IS Metrics for Traffic Engineering on page 180
List of Sample Output	<p>show isis interface on page 367</p> <p>show isis interface brief on page 367</p> <p>show isis interface detail on page 368</p> <p>show isis interface extensive on page 368</p>
Output Fields	<p>Table 17 on page 366 describes the output fields for the <code>show isis interface</code> command. Output fields are listed in the approximate order in which they appear.</p>

Table 17: 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 17: 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 routing devices 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 367](#).

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:
xe-6/1/0.0
  Index: 75, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding
  Adjacency advertisement: Advertise
  Level 1
    Adjacencies: 0, Priority: 64, Metric: 10
    Disabled
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 20.000 s, Hold Time: 60 s
    Designated Router: nemean.03
```

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 371
Output Fields	Table 18 on page 369 lists the output fields for the show isis overview command. Output fields are listed in the approximate order in which they appear.

Table 18: 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 18: 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.
Prefix export limit	Number of prefixes allowed to be exported, as configured by the prefix-export-limit statement.
Prefix export count	Number of prefixes exported.
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: 10.255.107.183
  Adjacency holddown: disabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled
  Traffic engineering: enabled
  Restart: Disabled
    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
  Prefix export limit: 5, Prefix export count: 5
  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 373</p> <p>show isis route (CLNS) on page 373</p> <p>show isis route on page 374</p>
Output Fields	<p>Table 19 on page 372 describes the output fields for the show isis route command. Output fields are listed in the approximate order in which they appear.</p>

Table 19: show isis route Output Fields

Field Name	Field Description
Current version	Number of the current version of the IS-IS routing table.

Table 19: 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 377 show isis spf results logical-system on page 378 show isis spf results (CLNS) on page 379
Output Fields	Table 20 on page 376 describes the output fields for the show isis spf command. Output fields are listed in the approximate order in which they appear.

Table 20: show isis spf Output Fields

Field Name	Field Description
Node	System ID of a node.
Metric	Metric to the node.

Table 20: 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
           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
           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 313
List of Sample Output	show isis statistics on page 383
Output Fields	Table 21 on page 382 describes the output fields for the show isis statistics command. Output fields are listed in the approximate order in which they appear.

Table 21: 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 policy

Syntax	show policy <logical-system (all <i>logical-system-name</i>)> < <i>policy-name</i> >
Syntax (EX Series Switches)	show policy < <i>policy-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 configured routing policies.
Options	<p>none—List the names of all configured routing policies.</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>policy-name</i>—(Optional) Show the contents of the specified policy.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show policy damping
List of Sample Output	show policy on page 384 show policy policy-name on page 385 show policy (Multicast Scoping) on page 385
Output Fields	Table 22 on page 384 lists the output fields for the show policy command. Output fields are listed in the approximate order in which they appear.

Table 22: show policy Output Fields

Field Name	Field Description
<i>policy-name</i>	Name of the policy listed.
<i>term</i>	Policy term listed.
<i>from</i>	Match condition for the policy.
<i>then</i>	Action for the policy.

Sample Output

show policy

```
user@host> show policy
```

```
Configured policies:
__vrf-export-red-internal__
__vrf-import-red-internal__
red-export
all_routes
```

show policy policy-name

```
user@host> show policy test-statics
Policy test-statics:
  from
    3.0.0.0/8  accept
    3.1.0.0/16  accept
  then reject
```

show policy (Multicast Scoping)

```
user@host> show policy test-statics
Policy test-statics:
  from
    multicast-scoping == 8
```

show policy conditions

Syntax	<pre>show policy conditions <condition-name> <detail> <dynamic> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switches)	<pre>show policy conditions <condition-name> <detail> <dynamic></pre>
Release Information	<p>Command introduced in Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Display all the configured conditions as well as the routing tables with which the configuration manager is interacting. If the detail keyword is included, the output also displays dependent routes for each condition.</p>
Options	<p>none—Display all configured conditions and associated routing tables.</p> <p>condition-name—(Optional) Display information about the specified condition only.</p> <p>detail—(Optional) Display the specified level of output.</p> <p>dynamic—(Optional) Display information about the conditions in the dynamic database.</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 policy conditions detail on page 387
Output Fields	<p>Table 23 on page 386 lists the output fields for the show policy conditions command. Output fields are listed in the approximate order in which they appear.</p>

Table 23: show policy conditions Output Fields

Field Name	Field Description	Level of Output
Condition	Name of configured condition.	All levels
event	Condition type. If the if-route-exists option is configured, the event type is: Existence of a route in a specific routing table.	All levels
Dependent routes	List of routes dependent on the condition, along with the latest generation number.	detail
Condition tables	List of routing tables associated with the condition, along with the latest generation number and number of dependencies.	All levels

Table 23: show policy conditions Output Fields (*continued*)

Field Name	Field Description	Level of Output
If-route-exists conditions	List of conditions configured to look for a route in the specified table.	All levels

Sample Output

show policy conditions detail

```
user@host> show policy conditions detail
Configured conditions:
Condition primary (static), event: Existence of a route in a specific routing
table
Dependent routes:
  8.41.0.0/24, generation 18

Condition standby (static), event: Existence of a route in a specific routing
table
Dependent routes:
  8.41.0.0/24, generation 18

Condition tables:
Table mpls.0, generation 0, dependencies 0, If-route-exists conditions: primary
(static) standby (static)
Table l3vpn.inet.0, generation 633, dependencies 2
```

show route

Syntax	show route <all> <destination-prefix> <logical-system (all <i>logical-system-name</i>)> <private>
Syntax (EX Series Switches)	show route <all> <destination-prefix> <private>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option private introduced in Junos OS Release 9.5. Option private introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the active entries in the routing tables.
Options	none —Display brief information about all active entries in the routing tables. all —(Optional) Display information about all routing tables, including private, or internal, routing tables. destination-prefix —(Optional) Display active entries for the specified 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. private —(Optional) Display information only about all private, or internal, routing tables.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring RIP</i>• <i>Example: Configuring RIPng</i>• Example: Configuring IS-IS on page 13• <i>Examples: Configuring Internal BGP Peering</i>• <i>Examples: Configuring External BGP Peering</i>• <i>Examples: Configuring OSPF Routing Policy</i>
List of Sample Output	show route on page 391 show route on page 391 show route destination-prefix on page 392 show route extensive on page 392

Output Fields Table 24 on page 389 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 24: 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.
[protocol, preference]	<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 24: 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. • Unverified—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • 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 24: 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). For VPNs, expect to see multiple push operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).

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

The following sample output shows a VPN route with composite next hops enabled. The first **Push** operation corresponds to the outer label. The second **Push** operation corresponds to the inner label.

```

user@host> show route 70.0.0.0

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0
hidden)

```

+ = Active Route, - = Last Active, * = Both

```
70.0.0.0/24      @[BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                  AS path: 13980 ?, validation-state: unverified
                  > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                  [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                  AS path: 13980 ?, validation-state: unverified
                  > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                  #[Multipath/255] 00:28:28, metric2 102
                  > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                  to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
```

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
    Validation State: unverified
    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	<p>none—Display all active routes.</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 active-path on page 393 show route active-path brief on page 394 show route active-path detail on page 394 show route active-path extensive on page 395 show route active-path terse on page 397
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 393](#).

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 advertising-protocol

Syntax	<code>show route advertising-protocol <i>protocol</i> <i>neighbor-address</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol.
Options	<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> <p><i>neighbor-address</i>—Address of the neighboring router to which the route entry is being transmitted.</p> <p><i>protocol</i>—Protocol transmitting the route:</p> <ul style="list-style-type: none"> • bgp—Border Gateway Protocol • dvmrp—Distance Vector Multicast Routing Protocol • msdp—Multicast Source Discovery Protocol • pim—Protocol Independent Multicast • rip—Routing Information Protocol • ripng—Routing Information Protocol next generation
Additional Information	Routes displayed are routes that the routing table has exported into the routing protocol and that have been filtered by the associated protocol's export routing policy statements.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the MED Attribute Directly</i>
List of Sample Output	show route advertising-protocol bgp (Layer 3 VPN) on page 400 show route advertising-protocol bgp detail on page 401 show route advertising-protocol bgp detail (Layer 2 VPN) on page 401 show route advertising-protocol bgp detail (Layer 3 VPN) on page 401 show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 401
Output Fields	Table 25 on page 399 lists the output fields for the show route advertising-protocol command. Output fields are listed in the approximate order in which they appear.

Table 25: show route advertising-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
<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) 	All levels
Prefix	Destination prefix.	brief none
<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
BGP group and type	BGP group name and type (Internal or External).	detail extensive
Route Distinguisher	Unique 64-bit prefix augmenting each IP subnet.	detail extensive
Advertised Label	Incoming label advertised by the LDP. When an IP packet enters a label-switched path (LSP), the ingress router examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label.	detail extensive
Label-Base, range	First label in a block of labels and label block size. A remote PE router uses this first label when sending traffic toward the advertising PE router.	detail extensive
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routers by advertising VPN labels. VPN labels transit over either an RSVP or an LDP LSP tunnel.	detail extensive
Nexthop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the next-hop advertisement to the peer is Self , and the RIB-out next hop is a specific IP address, the RIB-out IP address is included in the extensive output. See show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 401.	All levels
MED	Multiple exit discriminator value included in the route.	brief
Lclpref or Localpref	Local preference value included in the route.	All levels

Table 25: show route advertising-protocol Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 configured on the router, 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
Communities	Community path attribute for the route. See the output field table for 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 autonomous system (AS) that originated the route. These values are stored in the Attrset attribute at the originating router.	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 advertising-protocol bgp (Layer 3 VPN)

```

user@host> show route advertising-protocol bgp 10.255.14.171
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.172/32 Self              1      100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.181/32 Self              2      100 I

```

show route advertising-protocol bgp detail

```

user@host> show route advertising-protocol bgp 111.222.1.3 detail
bgp20.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
111.222.1.11/32 (1 entry, 1 announced)
  BGP group pe-pe type Internal
    Route Distinguisher: 111.255.14.11:69
    Advertised Label: 100000
    next hop: Self
    Localpref: 100
    AS path: 2 I
    Communities: target:69:20
    AIGP 210
111.8.0.0/16 (1 entry, 1 announced)
  BGP group pe-pe type Internal
    Route Distinguisher: 111.255.14.11:69
    Advertised Label: 100000
    Next hop: Self
    Localpref: 100
    AS path: 2 I
    Communities: target:69:20
    AIGP 210

```

show route advertising-protocol bgp detail (Layer 2 VPN)

```

user@host> show route advertising-protocol bgp 192.168.24.1 detail
vpn-a.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
192.168.16.1:1:1:1/96 (1 entry, 1 announced)
  BGP group int type Internal
    Route Distinguisher: 192.168.16.1:1
    Label-base : 32768, range : 3
    Nexthop: Self
    Localpref: 100
    AS path: I
    Communities: target:65412:100
    AIGP 210
    Layer2-info: encaps:VLAN, control flags:, mtu:

```

show route advertising-protocol bgp detail (Layer 3 VPN)

```

user@host> show route advertising-protocol bgp 10.255.14.176 detail
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
  BGP group ibgp type Internal
    Route Distinguisher: 10.255.14.174:2
    VPN Label: 101264
    Nexthop: Self
    Localpref: 100
    AS path: I
    Communities: target:200:100
    AIGP 210
    AttrSet AS: 100
      Localpref: 100
      AS path: I
  ...

```

show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address)

```

user@host> show route advertising-protocol bgp 200.0.0.2 170.0.1.0/24 extensive all
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 6 hidden)
  170.0.1.0/24 (2 entries, 1 announced)

```

```
BGP group eBGP-INTEROP type External
  Nexthop: Self (rib-out 10.100.3.2)
  AS path: [4713] 200 I
...
```


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 403
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 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 405 show route best detail on page 406 show route best extensive on page 407 show route best terse on page 407
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 406](#).

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	show route brief <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route brief <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 brief 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 brief on page 408
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 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. Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
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 419 show route detail (with BGP Multipath) on page 425 show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 425 show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 426
Output Fields	<p>Table 26 on page 410 describes the output fields for the show route detail command. Output fields are listed in the approximate order in which they appear.</p>

Table 26: 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 26: 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.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 27 on page 415 .

Table 26: 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 28 on page 416 .
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 26: 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	<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 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>
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. • Unverified—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.

Table 26: show route detail Output Fields (*continued*)

Field Name	Field Description
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
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 29 on page 418 for all possible values for this field.
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).
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 27 on page 415 describes all possible values for the Next-hop Types output field.

Table 27: 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.
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.

Table 27: Next-hop Types Output Field Values (*continued*)

Next-Hop Type	Description
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 28 on page 416 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 28: 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.

Table 28: State Output Field Values (*continued*)

Value	Description
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).
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.

Table 28: State Output Field Values (*continued*)

Value	Description
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 29 on page 418 describes the possible values for the Communities output field.

Table 29: 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.
bandwidth: local AS number:link-bandwidth-number	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.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	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.

Table 29: Communities Output Field Values (*continued*)

Value	Description
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	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.
route-type-vendor	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 area-number:ospf-route-type:options .
rte-type	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 area-number:ospf-route-type:options .
target	Defines which VPN the route participates in; target has the format 32-bit IP address:16-bit number . For example, 10.19.0.0:100.
unknown IANA	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.
unknown OSPF vendor community	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)

299840 (1 entry, 1 announced)

```

```

TSI:
KRT in-kerne 299840 /52 -> {indirect(1048575)}
    *RSVP Preference: 7/2
        Next hop type: Flood
        Address: 0x9174a30
        Next-hop reference count: 4
        Next hop type: Router, Next hop index: 798
        Address: 0x9174c28
        Next-hop reference count: 2
        Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
        Label-switched-path R2-to-R4-2p2mp
        Label operation: Pop
        Next hop type: Router, Next hop index: 1048574
        Address: 0x92544f0
        Next-hop reference count: 2
        Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
        Label-switched-path R2-to-R200-p2mp
        Label operation: Pop
        Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
        Label operation: Pop
        State: <Active Int>
        Age: 1:29 Metric: 1
        Task: RSVP
        Announcement bits (1): 0-KRT
        AS path: I...

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: l2 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
        Validation State: unverified
        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
        Validation State: unverified
        Task: BGP_2.10.1.1.6+58198
        AS path: 2 I
        Accepted MultipathContrib
        Localpref: 100
        Router ID: 1.1.1.3

```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
  *LDP   Preference: 9
        Next hop type: Flood
        Next-hop reference count: 3
        Address: 0x9097d90
        Next hop: via vt-0/1/0.1
        Next-hop index: 661
        Label operation: Pop
        Address: 0x9172130
        Next hop: via so-0/0/3.0
        Next-hop index: 654
        Label operation: Swap 299872
        State: **Active Int>
        Local AS: 1001

```

```

Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
    Next hop type: Flood
    Address: 0x2735208
    Next-hop reference count: 3
    Next hop type: Router, Next hop index: 1397
    Address: 0x2735d2c
    Next-hop reference count: 3
    Next hop: 1.3.8.2 via ge-1/2/22.0
    Label operation: Pop
    Load balance label: None;
    Next hop type: Router, Next hop index: 1395
    Address: 0x2736290
    Next-hop reference count: 3
    Next hop: 1.3.4.2 via ge-1/2/18.0
    Label operation: Pop
    Load balance label: None;
    State: <Active Int AckRequest MulticastRPF>
    Local AS: 10
    Age: 54:05      Metric: 1
    Validation State: unverified
    Task: LDP
    Announcement bits (1): 0-KRT
    AS path: I
    FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
    Primary Upstream : 1.1.1.3:0--1.1.1.2:0
      RPF Nexthops :
        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
    Backup Upstream : 1.1.1.3:0--1.1.1.6:0
      RPF Nexthops :
        ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
        ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```


show route exact

Syntax	<code>show route exact <i>destination-prefix</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route exact <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 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 427 show route exact detail on page 427 show route exact extensive on page 428 show route exact terse on page 428
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 430 show route export detail on page 430 show route export instance detail on page 430
Output Fields	Table 30 on page 429 lists the output fields for the show route export command. Output fields are listed in the approximate order in which they appear.

Table 30: show route export Output Fields

Field Name	Field Description	Level of Output
Table or <i>table-name</i>	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 30: 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                       Type: non-forwarding
Instance: red                         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 439 show route extensive (Access Route) on page 445 show route extensive (BGP PIC Edge) on page 446 show route extensive (FRR and LFA) on page 446 show route extensive (Route Reflector) on page 447 show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 447 show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 448
Output Fields	Table 31 on page 432 describes the output fields for the show route extensive command. Output fields are listed in the approximate order in which they appear.

Table 31: 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 31: 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). <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.
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 31: show route extensive Output Fields (*continued*)

Field Name	Field Description
PMSI	Provider multicast service interface (MVPN routing table).
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 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).
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.
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).

Table 31: show route extensive Output Fields (*continued*)

Field Name	Field Description
Indirect next hops	<p>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.</p> <p>When BGP PIC Edge is enabled, the output lines that contain Indirect next hop: weight follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values:</p> <ul style="list-style-type: none"> • 0x1 indicates active next hops. • 0x4000 indicates passive next hops.
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.
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Weight	<p>Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown.</p> <p>For sample output, see show route table.</p>

Table 31: show route extensive Output Fields (*continued*)

Field Name	Field Description
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.
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>

Table 31: show route extensive Output Fields (*continued*)

Field Name	Field Description
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>
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. • Unverified—Indicates that origin validation is not enabled for the BGP peers. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.

Table 31: show route extensive Output Fields (*continued*)

Field Name	Field Description
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
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.
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.

Table 31: show route extensive Output Fields (*continued*)

Field Name	Field Description
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-kerne1 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-kerne1 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-kerne1 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)
299840 (1 entry, 1 announced)
TSI:
KRT in-kernel 299840 /52 -> {indirect(1048575)}
  *RSVP Preference: 7/2
    Next hop type: Flood
    Address: 0x9174a30
    Next-hop reference count: 4
    Next hop type: Router, Next hop index: 798
    Address: 0x9174c28
    Next-hop reference count: 2
    Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
    Label-switched-path R2-to-R4-2p2mp

```

```

Label operation: Pop
Next hop type: Router, Next hop index: 1048574
Address: 0x92544f0
Next-hop reference count: 2
Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
Label-switched-path R2-to-R200-p2mp
Label operation: Pop
Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
Label operation: Pop
State: <Active Int>
Age: 1:29      Metric: 1
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I...

```

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.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
Label-base: 800008, range: 8, status-vector: 0x9F

```

```

...

12circuit.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: 12 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 (BGP PIC Edge)

```

user@host> show route 1.1.1.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  1.1.1.6/32 (3 entries, 2 announced)
    State: <CalcForwarding>
    TSI:
    KRT in-kerne1 1.1.1.6/32 -> {indirect(1048574), indirect(1048577)}
    Page 0 idx 0 Type 1 val 9219e30
    Nexthop: Self
    AS path: [2] 3 I
    Communities: target:2:1
    Path 1.1.1.6 from 1.1.1.4 Vector len 4. Val: 0
  ..
    #Multipath Preference: 255
    Next hop type: Indirect
    Address: 0x93f4010
    Next-hop reference count: 2
  ..
    Protocol next hop: 1.1.1.4
    Push 299824
    Indirect next hop: 944c000 1048574 INH Session ID: 0x3
    Indirect next hop: weight 0x1
    Protocol next hop: 1.1.1.5
    Push 299824
    Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
    Indirect next hop: weight 0x4000
    State: <ForwardingOnly Int Ext>
    Inactive reason: Forwarding use only
    Age: 25 Metric2: 15
    Validation State: unverified
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: 3 I
    Communities: target:2:1

```

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-kerne1 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 (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 label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
  *LDP Preference: 9
    Next hop type: Flood
    Next-hop reference count: 3
    Address: 0x9097d90
    Next hop: via vt-0/1/0.1
    Next-hop index: 661
    Label operation: Pop
    Address: 0x9172130
    Next hop: via so-0/0/3.0
    Next-hop index: 654
    Label operation: Swap 299872
    State: **Active Int>

```

```

Local AS: 1001
Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```

user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
    Next hop type: Flood
    Address: 0x2735208
    Next-hop reference count: 3
    Next hop type: Router, Next hop index: 1397
    Address: 0x2735d2c
    Next-hop reference count: 3
    Next hop: 1.3.8.2 via ge-1/2/22.0
    Label operation: Pop
    Load balance label: None;
    Next hop type: Router, Next hop index: 1395
    Address: 0x2736290
    Next-hop reference count: 3
    Next hop: 1.3.4.2 via ge-1/2/18.0
    Label operation: Pop
    Load balance label: None;
    State: <Active Int AckRequest MulticastRPF>
    Local AS: 10
    Age: 54:05      Metric: 1
    Validation State: unverified
    Task: LDP
    Announcement bits (1): 0-KRT
    AS path: I
    FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
      Primary Upstream : 1.1.1.3:0--1.1.1.2:0
        RPF Nexthops :
          ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
          ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
      Backup Upstream : 1.1.1.3:0--1.1.1.6:0
        RPF Nexthops :
          ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
          ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```

show route forwarding-table

Syntax	<pre> show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn> </pre>
Syntax (MX Series Routers)	<pre> show route forwarding-table <detail extensive summary> <all> <bridge-domain (all domain-name)> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <learning-vlan-id learning-vlan-id> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn> </pre>
Syntax (TX Matrix and TX Matrix Plus Routers)	<pre> show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <matching matching> <label name> <lcc number> <multicast> <table routing-instance-name> <vpn vpn> </pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Option bridge-domain introduced in Junos OS Release 7.5</p> <p>Option learning-vlan-id introduced in Junos OS Release 8.4</p> <p>Options all and vlan introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>

Description Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.



NOTE: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

Options **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

detail | extensive | summary—(Optional) Display the specified level of output.

all—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

bridge-domain (all | *bridge-domain-name*)—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

ccc *interface-name*—(Optional) Display route entries for the specified circuit cross-connect interface.

destination *destination-prefix*—(Optional) Destination prefix.

family *family*—(Optional) Display routing table entries for the specified family: **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

interface-name *interface-name*—(Optional) Display routing table entries for the specified interface.

label *name*—(Optional) Display route entries for the specified label.

lcc *number*—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

learning-vlan-id *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

table (default | *logical-system-name/routing-instance-name* | *routing-instance-name*)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the [show route instance](#) command.

vlan (all | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*—(Optional) Display routing table entries for a specified VPN.

Required Privilege Level view

List of Sample Output [show route forwarding-table on page 454](#)
[show route forwarding-table detail on page 455](#)
[show route forwarding-table destination extensive \(Weights and Balances\) on page 455](#)
[show route forwarding-table extensive on page 456](#)
[show route forwarding-table extensive \(RPF\) on page 457](#)
[show route forwarding-table family mpls on page 458](#)
[show route forwarding-table family vpls on page 458](#)
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled\) on page 458](#)
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled with MAC Statistics\) on page 459](#)
[show route forwarding-table family vpls extensive on page 459](#)
[show route forwarding-table table default on page 460](#)
[show route forwarding-table table logical-system-name/routing-instance-name on page 461](#)
[show route forwarding-table vpn on page 462](#)

Output Fields [Table 32 on page 452](#) lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 32: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the table logical-system-name/routing-instance-name option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels
Address family	Address family (for example, IP, IPv6, ISO, MPLS, and VPLS).	All levels
Destination	Destination of the route.	detail extensive
Route Type (Type)	<p>How the route was placed into the forwarding table. When the detail keyword is used, the route type might be abbreviated (as shown in parentheses):</p> <ul style="list-style-type: none"> • cloned (clon)—(TCP or multicast only) Cloned route. • destination (dest)—Remote addresses directly reachable through an interface. • destination down (iddn)—Destination route for which the interface is unreachable. • interface cloned (ifcl)—Cloned route for which the interface is unreachable. • route down (ifdn)—Interface route for which the interface is unreachable. • ignore (ignr)—Ignore this route. • interface (intf)—Installed as a result of configuring an interface. • permanent (perm)—Routes installed by the kernel when the routing table is initialized. • user—Routes installed by the routing protocol process or as a result of the configuration. 	All levels
Route Reference (RtRef)	Number of routes to reference.	detail extensive
Flags	<p>Route type flags:</p> <ul style="list-style-type: none"> • none—No flags are enabled. • accounting—Route has accounting enabled. • cached—Cache route. • incoming-iface interface-number—Check against incoming interface. • prefix load balance—Load balancing is enabled for this prefix. • rt nh decoupled—Route has been decoupled from the next hop to the destination. • sent to PFE—Route has been sent to the Packet Forwarding Engine. • static—Static route. 	extensive
Next hop	IP address of the next hop to the destination.	detail extensive

Table 32: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Next hop Type (Type)	<p>Next-hop type. When the detail keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> • broadcast (bcst)—Broadcast. • deny—Deny. • discard (dscd) —Discard. • hold—Next hop is waiting to be resolved into a unicast or multicast type. • indexed (idxd)—Indexed next hop. • indirect (indr)—Indirect next hop. • local (locl)—Local address on an interface. • routed multicast (mcrst)—Regular multicast next hop. • multicast (mcst)—Wire multicast next hop (limited to the LAN). • multicast discard (mdsc)—Multicast discard. • multicast group (mgrp)—Multicast group member. • receive (rcv)—Receive. • reject (rjct)—Discard. An ICMP unreachable message was sent. • resolve (rslv)—Resolving the 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. 	detail extensive
Index	Software index of the next hop that is used to route the traffic for a given prefix.	detail extensive none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	extensive
Reference (NhRef)	Number of routes that refer to this next hop.	detail extensive none
Next-hop interface (Netif)	Interface used to reach the next hop.	detail extensive 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 (see the Balance field description).	extensive
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	extensive
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when rpf-check is configured on the interface.	extensive

Sample Output

show route forwarding-table

```

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  46   4
0.0.0.0/32       perm  0                               dscd  44   1
1.1.1.0/24       ifdn  0                               rslv  608  1 ge-2/0/1.0
1.1.1.0/32       iddn  0 1.1.1.0           recv  606  1 ge-2/0/1.0
1.1.1.1/32       user  0                               rjct  46   4
1.1.1.1/32       intf  0 1.1.1.1           locl  607  2
1.1.1.1/32       iddn  0 1.1.1.1           locl  607  2
1.1.1.255/32     iddn  0 ff:ff:ff:ff:ff:ff bcst  605  1 ge-2/0/1.0
10.0.0.0/24      intf  0                               rslv  616  1 ge-2/0/0.0
10.0.0.0/32      dest  0 10.0.0.0          recv  614  1 ge-2/0/0.0
10.0.0.1/32      intf  0 10.0.0.1          locl  615  2
10.0.0.1/32      dest  0 10.0.0.1          locl  615  2
10.0.0.255/32    dest  0 10.0.0.255        bcst  613  1 ge-2/0/0.0
10.1.1.0/24      ifdn  0                               rslv  612  1 ge-2/0/1.0
10.1.1.0/32      iddn  0 10.1.1.0          recv  610  1 ge-2/0/1.0
10.1.1.1/32      user  0                               rjct  46   4
10.1.1.1/32      intf  0 10.1.1.1          locl  611  2
10.1.1.1/32      iddn  0 10.1.1.1          locl  611  2
10.1.1.255/32    iddn  0 ff:ff:ff:ff:ff:ff bcst  609  1 ge-2/0/1.0
10.209.0.0/16    user  0 10.209.63.254      ucst  419  20 fxp0.0
10.209.0.0/16    user  1 0:12:1e:ca:98:0    ucst  419  20 fxp0.0
10.209.0.0/18    intf  0                               rslv  418  1 fxp0.0
10.209.0.0/32    dest  0 10.209.0.0        recv  416  1 fxp0.0
10.209.2.131/32  intf  0 10.209.2.131      locl  417  2
10.209.2.131/32  dest  0 10.209.2.131      locl  417  2
10.209.17.55/32  dest  0 0:30:48:5b:78:d2  ucst  435  1 fxp0.0
10.209.63.42/32  dest  0 0:23:7d:58:92:ca  ucst  434  1 fxp0.0
10.209.63.254/32 dest  0 0:12:1e:ca:98:0    ucst  419  20 fxp0.0
10.209.63.255/32 dest  0 10.209.63.255     bcst  415  1 fxp0.0
10.227.0.0/16    user  0 10.209.63.254      ucst  419  20 fxp0.0

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  27   1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf  0                               locl  28   1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  6    1
ff00::/8         perm  0                               mdsc  4    1
ff02::1/128      perm  0 ff02::1           mcst  3    1

Routing table: ccc
MPLS:
Interface.Label  Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  16   1
100004(top)fe-0/0/1.0

```

show route forwarding-table detail

```

user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination      Type RtRef Next hop                Type Index NhRef Netif
default          user   2 0:90:69:8e:b1:1b        ucst  132   4 fxp0.0
default          perm   0                               rjct   14    1
10.1.1.0/24      intf   0 ff.3.0.21              ucst  322   1 so-5/3/0.0
10.1.1.0/32      dest   0 10.1.1.0                recv  324   1 so-5/3/0.0
10.1.1.1/32      intf   0 10.1.1.1                locl   321   1
10.1.1.255/32    dest   0 10.1.1.255              bcst  323   1 so-5/3/0.0
10.21.21.0/24    intf   0 ff.3.0.21              ucst  326   1 so-5/3/0.0
10.21.21.0/32    dest   0 10.21.21.0              recv  328   1 so-5/3/0.0
10.21.21.1/32    intf   0 10.21.21.1              locl   325   1
10.21.21.255/32  dest   0 10.21.21.255            bcst  327   1 so-5/3/0.0
127.0.0.1/32     intf   0 127.0.0.1              locl   320   1
172.17.28.19/32  clon    1 192.168.4.254          ucst  132   4 fxp0.0
172.17.28.44/32  clon    1 192.168.4.254          ucst  132   4 fxp0.0

...

Routing table: private1__inet
Internet:
Destination      Type RtRef Next hop                Type Index NhRef Netif
default          perm   0                               rjct   46    1
10.0.0.0/8       intf   0                               rslv  136   1 fxp1.0
10.0.0.0/32      dest   0 10.0.0.0                recv  134   1 fxp1.0
10.0.0.4/32      intf   0 10.0.0.4                locl   135   2
10.0.0.4/32      dest   0 10.0.0.4                locl   135   2

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop                Type Index NhRef Netif
default          perm   0                               rjct   38    1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop                Type Index NhRef Netif
default          perm   0                               rjct   22    1
ff00::/8         perm   0                               mdsc   21    1
ff02::1/128      perm   0 ff02::1                mcst   17    1

...

Routing table: mpls
MPLS:
Destination      Type RtRef Next hop                Type Index NhRef Netif
default          perm   0                               rjct  28    1

```

show route forwarding-table destination extensive (Weights and Balances)

```

user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination: 3.4.2.1/32
Route type: user
Route reference: 0                               Route interface-index: 0

```

Flags: sent to PFE		
Next-hop type: unicast	Index: 262143	Reference: 1
Nexthop: 4.4.4.4		
Next-hop type: unicast	Index: 335	Reference: 2
Next-hop interface: so-1/1/0.0	Weight: 22	Balance: 3
Nexthop: 145.12.1.2		
Next-hop type: unicast	Index: 337	Reference: 2
Next-hop interface: so-0/1/2.0	Weight: 33	Balance: 33

show route forwarding-table extensive

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:

Destination: default
  Route type: user
  Route reference: 2
  Flags: sent to PFE
  Nexthop: 0:90:69:8e:b1:1b
  Next-hop type: unicast
  Next-hop interface: fxp0.0
  Route interface-index: 0
  Index: 132      Reference: 4

Destination: default
  Route type: permanent
  Route reference: 0
  Flags: none
  Next-hop type: reject
  Route interface-index: 0
  Index: 14      Reference: 1

Destination: 127.0.0.1/32
  Route type: interface
  Route reference: 0
  Flags: sent to PFE
  Nexthop: 127.0.0.1
  Next-hop type: local
  Route interface-index: 0
  Index: 320      Reference: 1

...

Routing table: private1__inet [Index 1]
Internet:

Destination: default
  Route type: permanent
  Route reference: 0
  Flags: sent to PFE
  Next-hop type: reject
  Route interface-index: 0
  Index: 46      Reference: 1

Destination: 10.0.0.0/8
  Route type: interface
  Route reference: 0
  Flags: sent to PFE
  Next-hop type: resolve
  Next-hop interface: fxp1.0
  Route interface-index: 3
  Index: 136      Reference: 1

...

Routing table: iso [Index 0]
ISO:

Destination: default
  Route type: permanent

```

```

Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 38      Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 22      Reference: 1

Destination: ff00::/8
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: multicast discard
Route interface-index: 0
Index: 21      Reference: 1

...

Routing table: private1__inet6 [Index 1]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 54      Reference: 1

Destination: fe80::2a0:a5ff:fe3d:375/128
Route type: interface
Route reference: 0
Flags: sent to PFE
Next-hop: fe80::2a0:a5ff:fe3d:375
Next-hop type: local
Route interface-index: 0
Index: 75      Reference: 1

...

```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```

so-1/1/0 {
  unit 0 {
    family inet {
      rpf-check;
      address 15.95.1.2/30;
    }
  }
}

```

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination: 15.95.1.3/32
Route type: destination
Route reference: 0
Route interface-index: 67

```

```

Flags: sent to PFE
Nexthop: 15.95.1.3
Next-hop type: broadcast          Index: 328      Reference: 1
Next-hop interface: so-1/1/0.0
RPF interface: so-1/1/0.0

```

show route forwarding-table family mpls

```

user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0                user  0
1                user  0
2                user  0
100000           user  0 10.31.1.6          swap 100001      fe-1/1/0.0
800002           user  0                  Pop          vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                  user  0                  indr  351      4
                  Push 800000, Push 100002(top)

so-0/0/0.0

```

show route forwarding-table family vpls

```

user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          dymn  0
default          perm  0
fe-0/1/0.0       dymn  0
00:90:69:0c:20:1f/48      <<<<<Remote CE

                  dymn  0                  indr  351      4
                  Push 800000, Push 100002(top)

so-0/0/0.0
00:90:69:85:b0:1f/48      <<<<<Local CE

                  dymn  0                  ucst  354      2 fe-0/1/0.0

```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
lsi.1048832      intf  0
                  4.4.3.2          indr 1048574  4
                  Push 262145      621      2

ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0                  ucst  590      5 ge-2/3/9.0
0x30003/51       user  0                  comp  627      2
ge-2/3/9.0       intf  0                  ucst  590      5 ge-2/3/9.0
ge-3/1/3.0       intf  0                  ucst  619      4 ge-3/1/3.0
0x30002/51       user  0                  comp  600      2
0x30001/51       user  0                  comp  597      2

```


show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats
Destination      Type RtRef Next hop      Type Index  NhRef Netif
default          perm  0         4.4.3.2        dscd   519      1
1si.1048834      intf  0         4.4.3.2        indr   1048574  4
Push 262145      592      2
ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0         ucst   590      5 ge-2/3/9.0
0x30003/51       user  0         comp   630      2
ge-2/3/9.0       intf  0         ucst   590      5 ge-2/3/9.0
ge-3/1/3.0       intf  0         ucst   591      4 ge-3/1/3.0
0x30002/51       user  0         comp   627      2
0x30001/51       user  0         comp   624      2

```

show route forwarding-table family vpls extensive

```

user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

Destination: default
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Next-hop type: unicast
Next-hop interface: fe-0/1/2.0
Route interface-index: 72
Index: 289      Reference: 1
Index: 291      Reference: 3
Index: 290      Reference: 3

Destination: default
Route type: permanent
Route reference: 0
Flags: none
Next-hop type: discard
Route interface-index: 0
Index: 341      Reference: 1

Destination: fe-0/1/2.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: indirect
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Route interface-index: 69
Index: 293      Reference: 1
Index: 363      Reference: 4
Index: 301      Reference: 5
Index: 291      Reference: 3

Destination: fe-0/1/3.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Route interface-index: 70
Index: 292      Reference: 1

```

```

Next-hop type: indirect          Index: 363      Reference: 4
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: 10:00:00:01:01:01/48
Route type: dynamic
Route reference: 0               Route interface-index: 70
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0
Route used as destination:
  Packet count:      6640      Byte count:      675786
Route used as source
  Packet count:      6894      Byte count:      696424

Destination: 10:00:00:01:01:04/48
Route type: dynamic
Route reference: 0               Route interface-index: 69
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0
Route used as destination:
  Packet count:      96        Byte count:      8079
Route used as source:
  Packet count:      296        Byte count:      24955

Destination: 10:00:00:01:03:05/48
Route type: dynamic
Route reference: 0               Route interface-index: 74
Flags: sent to PFE, prefix load balance
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0

```

show route forwarding-table table default

```

user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0.0.0.0/32       perm  0
10.0.60.0/30     user  0 10.0.60.13          ucst  713  5 fe-0/1/3.0
10.0.60.12/30    intf  0
10.0.60.12/32    dest  0 10.0.60.12          recv  686  1 fe-0/1/3.0
10.0.60.13/32    dest  0 0:5:85:8b:bc:22     ucst  713  5 fe-0/1/3.0
10.0.60.14/32    intf  0 10.0.60.14          locl  687  2
10.0.60.14/32    dest  0 10.0.60.14          locl  687  2
10.0.60.15/32    dest  0 10.0.60.15          bcst  685  1 fe-0/1/3.0
10.0.67.12/30    user  0 10.0.60.13          ucst  713  5 fe-0/1/3.0
10.0.80.0/30     ifdn  0 ff.3.0.21          ucst  676  1 so-0/0/1.0
10.0.80.0/32     dest  0 10.0.80.0          recv  678  1 so-0/0/1.0
10.0.80.2/32     user  0
10.0.80.2/32     intf  0 10.0.80.2          locl  675  1

```

```

10.0.80.3/32      dest    0 10.0.80.3      bcst  677    1 so-0/0/1.0
10.0.90.12/30     intf    0                rslv  684    1 fe-0/1/0.0
10.0.90.12/32     dest    0 10.0.90.12   recv  682    1 fe-0/1/0.0
10.0.90.14/32     intf    0 10.0.90.14   locl  683    2
10.0.90.14/32     dest    0 10.0.90.14   locl  683    2
10.0.90.15/32     dest    0 10.0.90.15   bcst  681    1 fe-0/1/0.0
10.5.0.0/16       user    0 192.168.187.126 ucst  324   15 fxp0.0
10.10.0.0/16      user    0 192.168.187.126 ucst  324   15 fxp0.0
10.13.10.0/23     user    0 192.168.187.126 ucst  324   15 fxp0.0
10.84.0.0/16      user    0 192.168.187.126 ucst  324   15 fxp0.0
10.150.0.0/16     user    0 192.168.187.126 ucst  324   15 fxp0.0
10.157.64.0/19    user    0 192.168.187.126 ucst  324   15 fxp0.0
10.209.0.0/16     user    0 192.168.187.126 ucst  324   15 fxp0.0

```

...

Routing table: default.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	60	1	

Routing table: default.inet6

Internet6:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	44	1	
::/128	perm	0		dscd	42	1	
ff00::/8	perm	0		mdsc	43	1	
ff02::1/128	perm	0	ff02::1	mcst	39	1	

Routing table: default.mpls

MPLS:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	50	1	

show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
```

Logical system: R4

Routing table: vpn-red.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	563	1	
0.0.0.0/32	perm	0		dscd	561	2	
1.0.0.1/32	user	0		dscd	561	2	
2.0.2.0/24	intf	0		rslv	771	1	ge-1/2/0.3
2.0.2.0/32	dest	0	2.0.2.0	recv	769	1	ge-1/2/0.3
2.0.2.1/32	intf	0	2.0.2.1	locl	770	2	
2.0.2.1/32	dest	0	2.0.2.1	locl	770	2	
2.0.2.2/32	dest	0	0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0	ucst	789	1	ge-1/2/0.3
2.0.2.255/32	dest	0	2.0.2.255	bcst	768	1	ge-1/2/0.3
224.0.0.0/4	perm	1		mdsc	562	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	558	1	
255.255.255.255/32	perm	0		bcst	559	1	

Logical system: R4

Routing table: vpn-red.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	608	1	

```

Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct  708    1
::/128           perm  0                dscd  706    1
ff00::/8         perm  0                mdsc  707    1
ff02::1/128      perm  0 ff02::1          mcst  704    1

```

```

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                dscd  638

```

show route forwarding-table vpn

```

user@host> show route forwarding-table vpn VPN-A
Routing table:: VPN-A.inet
Internet:
Destination      Type RtRef Nexthop          Type Index NhRef Netif
default          perm  0                rjct   4    4
10.39.10.20/30   intf  0 ff.3.0.21          ucst   40    1
so-0/0/0.0
10.39.10.21/32   intf  0 10.39.10.21        locl   36    1
10.255.14.172/32 user   0                ucst   69    2
so-0/0/0.0
10.255.14.175/32 user   0                indr   81    3
Push 100004, Push
100004(top) so-1/0/0.0
224.0.0.0/4      perm  2                mdsc   5    3
224.0.0.1/32     perm  0 224.0.0.1          mcst   1    8
224.0.0.5/32     user   1 224.0.0.5          mcst   1    8
255.255.255.255/32 perm  0                bcst   2    3

```

show route hidden

Syntax	show route hidden <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display only hidden route information. A hidden route is unusable, even if it is the best path.
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>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Hidden Routes</i>
List of Sample Output	show route hidden on page 463 show route hidden detail on page 464 show route hidden extensive on page 464 show route hidden terse on page 464
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 hidden

```

user@host> show route hidden
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
127.0.0.1/32      [Direct/0] 04:26:38
                  > via lo0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.5.5.5/32      [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                  AS path: 100 I
                  Unusable
10.12.1.0/24     [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                  AS path: 100 I
                  Unusable

```

```

10.12.80.4/30      [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                  AS path: I
                  Unusable
...

```

show route hidden detail

```

user@host> show route hidden detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
127.0.0.1/32 (1 entry, 0 announced)
    Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Hidden Martian Int>
        Local AS:      1
        Age: 4:27:37
        Task: IF
        AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.5.5.5/32 (1 entry, 0 announced)
    BGP Preference: 170/-101
        Route Distinguisher: 10.4.4.4:4
        Next hop type: Unusable
        Next-hop reference count: 6
        State: <Secondary Hidden Int Ext>
        Local AS:      1 Peer AS:      1
        Age: 3:45:09
        Task: BGP_1.10.4.4.4+2493
        AS path: 100 I
        Communities: target:1:999
        VPN Label: 100064
        Localpref: 100
        Router ID: 10.4.4.4
        Primary Routing Table bgp.13vpn.0

...

```

show route hidden extensive

The output for the **show route hidden extensive** command is identical to that of the **show route hidden detail** command. For sample output, see [show route hidden detail on page 464](#).

show route hidden terse

```

user@host> show route hidden 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
127.0.0.1/32      D  0                >100.0

```

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

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
10.5.5.5/32	B 170	100		Unusable	100 I
10.12.1.0/24	B 170	100		Unusable	100 I
10.12.80.4/30	B 170	100		Unusable	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

+ = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
10.4.4.4:4:10.5.5.5/32	B 170	100		Unusable	100 I
10.4.4.4:4:10.12.1.0/24	B 170	100		Unusable	100 I
10.4.4.4:4:10.12.80.4/30	B 170	100		Unusable	I

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

Syntax	show route inactive-path <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route inactive-path <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 routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.
Options	<p>none—Display all inactive routes.</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-path on page 466 show route inactive-path detail on page 467 show route inactive-path extensive on page 468 show route inactive-path terse on page 468
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 467](#).

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 instance

Syntax	<pre>show route instance <brief detail summary> <instance-name> <logical-system (all logical-system-name)> <operational></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show route instance <brief detail summary> <instance-name> <operational></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display 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 logical-system-name)—(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
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling</i> • <i>Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart</i>
List of Sample Output	<p>show route instance on page 470</p> <p>show route instance detail (Graceful Restart Complete) on page 471</p> <p>show route instance detail (Graceful Restart Incomplete) on page 473</p> <p>show route instance detail (VPLS Routing Instance) on page 474</p> <p>show route instance operational on page 475</p> <p>show route instance summary on page 475</p>
Output Fields	<p>Table 33 on page 470 lists the output fields for the show route instance command. Output fields are listed in the approximate order in which they appear.</p>

Table 33: 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
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
Active/holddown/hidden

```

```

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 next-hop

Syntax	<code>show route next-hop <i>next-hop</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 next-hop <i>next-hop</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 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 476 show route next-hop detail on page 477 show route next-hop extensive on page 479 show route next-hop terse on page 480
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 output

Syntax	show route output (address <i>ip-address</i> interface <i>interface-name</i>) <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route output (address <i>ip-address</i> interface <i>interface-name</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	<p>Display the entries in the routing table learned through static routes and interior gateway protocols that are to be sent out the interface with either the specified IP address or specified name.</p> <p>To view routes advertised to a neighbor or received from a neighbor for the BGP protocol, use the show route advertising-protocol bgp and show route receive-protocol bgp commands instead.</p>
Options	<p>address <i>ip-address</i>—Display entries in the routing table that are to be sent out the interface with the specified IP address.</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>interface <i>interface-name</i>—Display entries in the routing table that are to be sent out the interface with the specified name.</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 output address on page 483 show route output address detail on page 483 show route output address extensive on page 484 show route output address terse on page 484 show route output interface on page 484 show route output interface detail on page 485 show route output interface extensive on page 485 show route output interface terse on page 485
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 output address

```

user@host> show route output address 36.1.1.1/24

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

36.1.1.0/24          *[Direct/0] 00:19:56
                    > via so-0/1/2.0
                    [OSPF/10] 00:19:55, metric 1
                    > via so-0/1/2.0

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: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route output address detail

```

user@host> show route output address 36.1.1.1 detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
36.1.1.0/24 (2 entries, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via so-0/1/2.0, selected
    State: <Active Int>
    Age: 23:00
    Task: IF
    AS path: I
  OSPF Preference: 10
    Next-hop reference count: 1
    Next hop: via so-0/1/2.0, selected
    State: <Int>
    Inactive reason: Route Preference
    Age: 22:59      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)

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: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route output address extensive

The output for the **show route output address extensive** command is identical to that of the **show route output address detail** command. For sample output, see [show route output address detail on page 483](#).

show route output address terse

```
user@host> show route output address 36.1.1.1 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
* 36.1.1.0/24      D   0                >so-0/1/2.0
                   O  10             1         >so-0/1/2.0

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: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output interface

```
user@host> show route output interface so-0/1/2.0

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.240/32  *[OSPF/10] 00:13:00, metric 2
                  via so-0/1/2.0
                  > via so-0/3/2.0
10.255.71.241/32  *[OSPF/10] 00:13:10, metric 1
                  > via so-0/1/2.0
14.1.1.0/24       *[OSPF/10] 00:05:11, 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] 00:13:10, metric 2
                  > via so-0/1/2.0
36.1.1.0/24       *[Direct/0] 00:13:21
                  > via so-0/1/2.0
                  [OSPF/10] 00:13:20, metric 1
                  > via so-0/1/2.0

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: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output interface detail

```
user@host> show route output interface so-0/1/2.0 detail
```

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.255.71.240/32 (1 entry, 1 announced)
    *OSPF    Preference: 10
              Next-hop reference count: 2
              Next hop: via so-0/1/2.0
              Next hop: via so-0/3/2.0, selected
              State: <Active Int>
              Age: 14:52      Metric: 2
              Area: 0.0.0.0
              Task: OSPF
              Announcement bits (1): 0-KRT
              AS path: I

10.255.71.241/32 (1 entry, 1 announced)
    *OSPF    Preference: 10
              Next-hop reference count: 4
              Next hop: via so-0/1/2.0, selected
              State: <Active Int>
              Age: 15:02      Metric: 1
              Area: 0.0.0.0
              Task: OSPF
              Announcement bits (1): 0-KRT
              AS path: I

...
```

show route output interface extensive

The output for the **show route output interface extensive** command is identical to that of the **show route output interface detail** command. For sample output, see [show route output interface detail on page 485](#).

show route output interface terse

```
user@host> show route output interface so-0/1/2.0 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
* 10.255.71.240/32  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
* 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
* 36.1.1.0/24       D  0
                        0 10    1          >so-0/1/2.0
                        >so-0/1/2.0

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: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

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 489 show route protocol access-internal extensive on page 489 show route protocol arp on page 489 show route protocol bgp on page 490 show route protocol bgp detail on page 490 show route protocol bgp extensive on page 490 show route protocol bgp terse on page 491 show route protocol direct on page 491 show route protocol frr on page 492 show route protocol l2circuit detail on page 492 show route protocol l2vpn extensive on page 493 show route protocol ldp on page 494 show route protocol ldp extensive on page 494 show route protocol ospf (Layer 3 VPN) on page 495 show route protocol ospf detail on page 496 show route protocol rip on page 496 show route protocol rip detail on page 496 show route protocol ripng table inet6 on page 497 show route protocol static detail on page 497

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 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 502 show route receive-protocol bgp extensive on page 502 show route receive-protocol bgp table extensive on page 502 show route receive-protocol bgp logical-system extensive on page 503 show route receive-protocol bgp detail (Layer 2 VPN) on page 504 show route receive-protocol bgp extensive (Layer 2 VPN) on page 504 show route receive-protocol bgp (Layer 3 VPN) on page 505 show route receive-protocol bgp detail (Layer 3 VPN) on page 505 show route receive-protocol bgp extensive (Layer 3 VPN) on page 506	
Output Fields	Table 34 on page 499 describes the output fields for the show route receive-protocol command. Output fields are listed in the approximate order in which they appear.	

Table 34: 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 34: show route receive-protocol Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>number</i> 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
<i>destination-prefix</i> (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 34: 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      Lclpref AS path
inet.3: 4 destinations, 4 routes (4 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: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix          Nexthop          MED      Lclpref 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      Lclpref 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      Lclpref AS path
inet.3: 4 destinations, 4 routes (4 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: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref 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      Lclpref 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 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	<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> <p><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).</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show route summary
List of Sample Output	show route table bgp.l2.vpn on page 508 show route table bgp.l3vpn.0 on page 508 show route table bgp.l3vpn.0 detail on page 508 show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 510 show route table inet.0 on page 510 show route table inet6.0 on page 510 show route table inet6.3 on page 511 show route table inetflow detail on page 511 show route table l2circuit.0 on page 511 show route table mpls on page 512 show route table mpls extensive on page 512 show route table mpls.0 on page 512 show route table mpls.0 detail (PTX Series) on page 513 show route table mpls.0 extensive (PTX Series) on page 513 show route table mpls.0 (RSVP Route—Transit LSP) on page 514 show route table vpls_1 detail on page 514 show route table vpn-a on page 515 show route table vpn-a.mdt.0 on page 515 show route table VPN-A detail on page 515 show route table VPN-AB.inet.0 on page 516

[show route table VPN_blue.mvpn-inet6.0 on page 516](#)

[show route table vrf1.mvpn.0 extensive on page 517](#)

[show route table inetflow detail on page 517](#)

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.l2.vpn

```
user@host> show route table bgp.l2.vpn
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 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 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 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 detail (PTX Series)

```

user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
  *L2VPN Preference: 7
    Next hop type: Indirect
    Address: 0x9438f34
    Next-hop reference count: 2
    Next hop type: Router, Next hop index: 567
    Next hop: 3.0.0.1 via ge-0/0/1.0, selected
    Label operation: Push 299808
    Label TTL action: prop-ttl
    Load balance label: Label 299808:None;
    Session Id: 0x1
    Protocol next hop: 10.255.255.1
    Label operation: Push 299872 Offset: 252
    Label TTL action: no-prop-ttl
    Load balance label: Label 299872:Flow label PUSH;
    Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
    Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
    State: <Active Int>
    Age: 21 Metric2: 1
    Validation State: unverified
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 2-Common L2 VC
    AS path: I

```

show route table mpls.0 extensive (PTX Series)

```

user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0 /32 -> {composite(570)}
  *L2VPN Preference: 7
    Next hop type: Indirect
    Address: 0x9438f34
    Next-hop reference count: 2
    Next hop type: Router, Next hop index: 567
    Next hop: 3.0.0.1 via ge-0/0/1.0, selected
    Label operation: Push 299808
    Label TTL action: prop-ttl
    Load balance label: Label 299808:None;
    Session Id: 0x1
    Protocol next hop: 10.255.255.1
    Label operation: Push 299872 Offset: 252
    Label TTL action: no-prop-ttl
    Load balance label: Label 299872:Flow label PUSH;
    Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
    Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
    State: <Active Int>
    Age: 47 Metric2: 1
    Validation State: unverified
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 2-Common L2 VC

```

```

AS path: I
Composite next hops: 1
  Protocol next hop: 10.255.255.1 Metric: 1
  Label operation: Push 299872 Offset: 252
  Label TTL action: no-prop-ttl
  Load balance label: Label 299872:Flow label PUSH;
  Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
  Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 3.0.0.1 via ge-0/0/1.0
    Session Id: 0x1
  10.255.255.1/32 Originating RIB: inet.3
  Metric: 1 Node path count: 1
  Forwarding nexthops: 1
  Nexthop: 3.0.0.1 via ge-0/0/1.0

```

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.0.0.0.0+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.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
    *[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 vrf1.mvpn.0 extensive

```

user@host> show route table vrf1.mvpn.0 extensive
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
    *MVPN    Preference: 70
              PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
    Next hop type: Indirect
    Address: 0xbb2c944
    Next-hop reference count: 360
    Protocol next hop: 10.255.50.77
    Indirect next hop: 0x0 - INH Session ID: 0x0
    State: <Active Int Ext>
    Age: 53:03      Metric2: 1
    Validation State: unverified
    Task: mvpn global task
    Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export

    AS path: I

```

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

user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
hidden)
22.0.0.0/32 (3 entries, 1 announced)
    State: <OnList CalcForwarding>
TSI:
KRT in-kernel 22.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 22.0.0.0 from 2.3.0.0 Vector len 4. Val: 1
    @BGP Preference: 170/-1
    Route Distinguisher: 2:1
    Next hop type: Indirect
    Address: 0x258059e4
    Next-hop reference count: 2
    Source: 2.2.0.0
    Next hop type: Router
    Next hop: 10.1.1.1 via ge-1/1/9.0, selected
    Label operation: Push 707633
    Label TTL action: prop-ttl
    Session Id: 0x17d8
    Protocol next hop: 2.2.0.0
    Push 16
    Composite next hop: 0x25805988 - INH Session ID: 0x193c
    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
    State: <Secondary Active Int Ext ProtectionPath ProtectionCand>
    Local AS: 2 Peer AS: 2
    Age: 23 Metric2: 35
    Validation State: unverified
    Task: BGP_2.2.2.0.0+34549
    AS path: I
    Communities: target:2:1
    Import Accepted
    VPN Label: 16
    Localpref: 0
    Router ID: 2.2.0.0
    Primary Routing Table bgp.13vpn.0
    Composite next hops: 1
        Protocol next hop: 2.2.0.0 Metric: 35
        Push 16
        Composite next hop: 0x25805988 - INH Session ID: 0x193c
        Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
        Indirect path forwarding next hops: 1
            Next hop type: Router
            Next hop: 10.1.1.1 via ge-1/1/9.0

```

```

                Session Id: 0x17d8
                2.2.0.0/32 Originating RIB: inet.3
                Metric: 35                      Node path count: 1
                Forwarding nexthops: 1
                Nexthop: 10.1.1.1 via ge-1/1/9.0
BGP  Preference: 170/-1
      Route Distinguisher: 2:1
      Next hop type: Indirect
      Address: 0x9347028
      Next-hop reference count: 3
      Source: 2.3.0.0
      Next hop type: Router, Next hop index: 702
      Next hop: 10.1.4.2 via ge-1/0/0.0, selected
      Label operation: Push 634278
      Label TTL action: prop-ttl
      Session Id: 0x17d9
      Protocol next hop: 2.3.0.0
      Push 16
      Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
      Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
      State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

      Inactive reason: Not Best in its group - IGP metric
      Local AS:      2 Peer AS:      2
      Age: 3:34      Metric2: 70
      Validation State: unverified
      Task: BGP_2.2.3.0.0+32805
      Announcement bits (2): 0-KRT 1-BGP_RT_Background
      AS path: I
      Communities: target:2:1
      Import Accepted
      VPN Label: 16
      Localpref: 0
      Router ID: 2.3.0.0
      Primary Routing Table bgp.13vpn.0
      Composite next hops: 1
        Protocol next hop: 2.3.0.0 Metric: 70
        Push 16
        Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
        Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da
        Indirect path forwarding next hops: 1
          Next hop type: Router
          Next hop: 10.1.4.2 via ge-1/0/0.0
          Session Id: 0x17d9
          2.3.0.0/32 Originating RIB: inet.3
          Metric: 70                      Node path count: 1
          Forwarding nexthops: 1
          Nexthop: 10.1.4.2 via ge-1/0/0.0
#Multipath Preference: 255
  Next hop type: Indirect
  Address: 0x24afca30
  Next-hop reference count: 1
  Next hop type: Router
  Next hop: 10.1.1.1 via ge-1/1/9.0, selected
  Label operation: Push 707633
  Label TTL action: prop-ttl
  Session Id: 0x17d8
  Next hop type: Router, Next hop index: 702
  Next hop: 10.1.4.2 via ge-1/0/0.0

```

Label operation: Push 634278
Label TTL action: prop-ttl
Session Id: 0x17d9
Protocol next hop: 2.2.0.0
Push 16
Composite next hop: 0x25805988 - INH Session ID: 0x193c
Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

Protocol next hop: 2.3.0.0
Push 16
Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight

0x4000

State: <ForwardingOnly Int Ext>
Inactive reason: Forwarding use only
Age: 23 Metric2: 35
Validation State: unverified
Task: RT
AS path: I
Communities: target:2:1

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 523
Output Fields	Table 35 on page 521 describes the output fields for the show route terse command. Output fields are listed in the approximate order in which they appear.

Table 35: 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 35: show route terse Output Fields (*continued*)

Field Name	Field Description
<i>route key</i>	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 35: 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      I
?                               B 170      100           >10.0.0.2      I
unverified
* ? 1.1.1.1/32       D 0           >10.0.0.2      200 I
* V 2.2.0.2/32       B 170     110         >10.0.0.2
valid
* ? 10.0.0.0/30      D 0           >1t-1/2/0.1    I
?                               B 170     100           >10.0.0.2
unverified
* ? 10.0.0.1/32      L 0           Local          I
* ? 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

```

show security keychain

Syntax	show security keychain <brief detail>
Release Information	Command introduced in Junos OS Release 11.2.
Description	Display information about authentication keychains configured 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.
Options	none —Display information about authentication keychains. brief detail —(Optional) Display the specified level of output.
Required Privilege Level	view
List of Sample Output	show security keychain brief on page 526 show security keychain detail on page 526
Output Fields	Table 10 on page 339 describes the output fields for the show security keychain command. Output fields are listed in the approximate order in which they appear.

Table 36: show security keychain Output Fields

Field Name	Field Description	Level of Output
keychain	The name of the keychain in operation.	All levels
Active-ID Send	Number of routing protocols packets sent with the active key.	All levels
Active-ID Receive	Number of routing protocols packets received with the active key.	All levels
Next-ID Send	Number of routing protocols packets sent with the next key.	All levels
Next-ID Receive	Number of routing protocols packets received with the next key.	All levels
Transition	Amount of time until the current key will be replaced with the next key in the keychain.	All levels
Tolerance	Configured clock-skew tolerance, in seconds, for accepting keys for a key chain.	All levels
Id	Identification number configured for the current key.	detail

Table 36: show security keychain Output Fields (*continued*)

Field Name	Field Description	Level of Output
Algorithm	Authentication algorithm configured for the current key.	detail
State	<p>State of the current key.</p> <p>The value can be:</p> <ul style="list-style-type: none"> • receive • send • send-receive <p>For the active key, the State can be send-receive, send, or receive. For keys that have a future start time, the State is inactive. Compare the State field to the Mode field.</p>	detail
Option	<p>For IS-IS only, the option determines how Junos OS encodes the message authentication code in routing protocol packets.</p> <p>The values can be:</p> <ul style="list-style-type: none"> • basic—Based on RFC 5304. • isis-enhanced—Based on RFC 5310. <p>The default value is basic. When you configure the isis-enhanced option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.</p> <p>When you configure basic (or do not include the options statement in the key configuration) Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices.</p> <p>Because this setting is for IS-IS only, the TCP and the BFD protocol ignore the encoding option configured in the key.</p>	detail
Start-time	Time that the current key became active.	detail

Table 36: show security keychain Output Fields (*continued*)

Field Name	Field Description	Level of Output
Mode	<p>Mode of each key (Informational only.)</p> <p>The value can be</p> <ul style="list-style-type: none"> • receive • send • send-receive <p>The mode of the key is based on the configuration. Suppose you configure two keys, one with a start-time of today and the other with a start-time of next week. For both keys, the Mode can be send-receive, send, or receive, regardless of the configured start-time. Compare the Mode field to the State field.</p>	detail

Sample Output

show security keychain brief

```

user@host> show security keychain brief
keychain              Active-ID      Next-ID      Transition  Tolerance
                     Send  Receive    Send  Receive
hakr                   3     3           1     1         1d 23:58    3600

```


show security keychain detail

```

user@host> show security keychain detail
keychain              Active-ID      Next-ID      Transition  Tolerance
                     Send  Receive    Send  Receive
hakr                   3     3           1     1         1d 23:58    3600
  Id 3, Algorithm hmac-md5, State send-receive, Option basic
  Start-time Wed Aug 11 16:28:00 2010, Mode send-receive
  Id 1, Algorithm hmac-md5, State inactive, Option basic
  Start-time Fri Aug 20 11:30:57 2010, Mode send-receive

```

test policy

Syntax	<code>test policy <i>policy-name</i> <i>prefix</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	Test a policy configuration to determine which prefixes match routes in the routing table.
<div>  <p>NOTE: If you are using the <code>test policy</code> command on a logical system, you must first set the CLI to the logical system context. For example, if you want to test a routing policy that is configured on logical system R2, first run the <code>set cli logical-system R2</code> command.</p> </div>	
Options	<p><i>policy-name</i>—Name of a policy.</p> <p><i>prefix</i>—Destination prefix to match.</p>
Additional Information	All prefixes in the default unicast routing table (inet.0) that match prefixes that are the same as or longer than the specific prefix are processed by the from clause in the specified policy. All prefixes accepted by the policy are displayed. The test policy command evaluates a policy differently from the BGP import process. When testing a policy that contains an interface match condition in the from clause, the test policy command uses the match condition. In contrast, BGP does not use the interface match condition when evaluating the policy against routes learned from internal BGP (IBGP) or external BGP (EGBP) multihop peers.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Routing Policy Tests</i> • <i>Example: Testing a Routing Policy with Complex Regular Expressions</i>
List of Sample Output	test policy on page 527
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

test policy

```
user@host> test policy test-statics 3.0.0.1/8
inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
Prefixes passing policy:
```

```
3.0.0.0/8      *[BGP/170] 16:22:46, localpref 100, from 10.255.255.41
               AS Path: 50888 I
               > to 10.11.4.32 via en0.2, label-switched-path 12
3.3.3.1/32     *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
3.3.3.2/32     *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
3.3.3.3/32     *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
3.3.3.4/32     *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
Policy test-statics: 5 prefixes accepted, 0 prefixes rejected
```

traceroute clns

Syntax	traceroute clns <host> routing-instance <name> source <source-address> ttl <value> wait <seconds>
Release Information	Command introduced in Junos OS Release 8.0
Description	Trace the route belonging to Connectionless Network Service (CLNS).
Options	<p>host—IP address or name of remote host.</p> <p>routing-instance name—Name of the routing instance for a traceroute attempt.</p> <p>source source address—Source address to be used in outgoing packets.</p> <p>ttl value—CLNP maximum time-to-live value. The range of values is 1 through 255.</p> <p>wait seconds—Number of seconds to wait for a response. The range of values is 1 second through 1 day.</p>
Required Privilege Level	network
List of Sample Output	traceroute clns on page 529
Output Fields	Table 37 on page 529 describes the output fields for the traceroute clns command. Output fields are listed in the approximate order in which they appear.

Table 37: Traceroute clns Output Fields

Field Name	Field Description
source	Source address used in outgoing packets.
traceroute to	IP address of the receiver.
hops max	Maximum number of hops allowed.
byte packets	Size of packets being sent.

Sample Output

traceroute clns

```

user@host>traceroute clns <ISO address of the destination> source <ISO address of the source>
traceroute clns 49.0005.80ff.f800.0000.0108.0001.0102.5522.4145.00 source
47.0005.80ff.f800.0000.0108.0001.0102.5522.4143.00
clnstraceroute to 49.0005.80ff.f800.0000.0108.0001.0102.5522.4145.00 from
47.0005.80ff.f800.0000.0108.0001.0102.5522.4143.00, 30 hops max, 17 byte packets

```

```
1 47.0005.80ff.f800.0000.0108.0001.0010.0010.0010.00 7.080 ms 5.579 ms 5.882
ms
2 * * *
3 * * *
4 * * *
5 49.0005.80ff.f800.0000.0108.0001.0102.5522.4145.00 11.325 ms 7.704 ms 5.261
ms
```


PART 4

Troubleshooting

- [IS-IS Troubleshooting on page 533](#)
- [Routing Protocol Process Memory FAQs on page 555](#)

IS-IS Troubleshooting

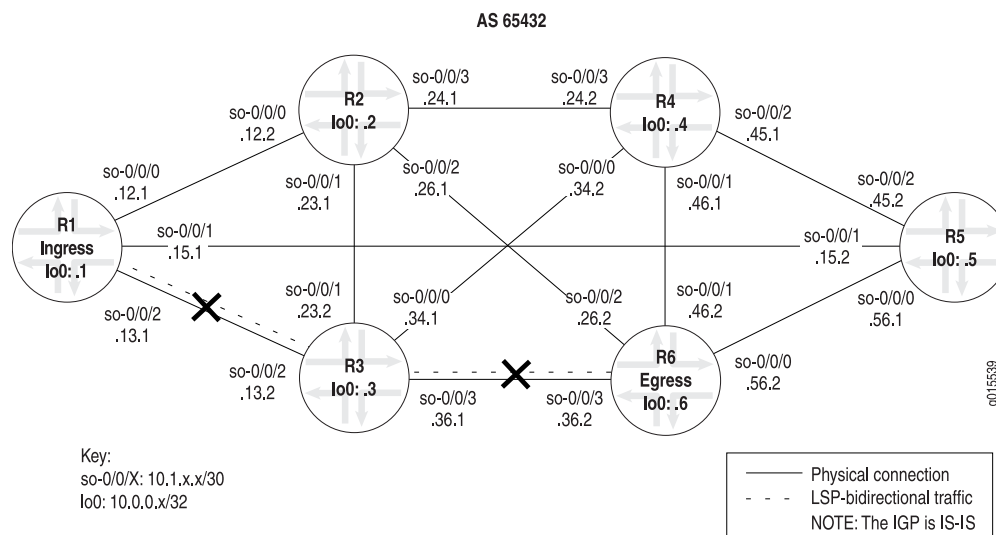
- [Verifying the IS-IS Protocol on page 533](#)
- [Verifying the IS-IS Configuration on a Router in a Network on page 541](#)
- [Displaying the Status of IS-IS Adjacencies on page 547](#)
- [Displaying Detailed IS-IS Protocol Information on page 548](#)
- [Analyzing IS-IS Link-State PDUs in Detail on page 550](#)
- [Displaying Sent or Received IS-IS Protocol Packets on page 552](#)

Verifying the IS-IS Protocol

Purpose If your MPLS network is configured with IS-IS as the interior gateway protocol (IGP), and the output of the **show mpls lsp extensive** command shows that there is a problem, check the IP and IS-IS layers. Because IS-IS and IP are independent of each other, you can check either layer first. For more information about checking the IP layer, see *Verifying the IP Layer*.

After you have checked the IP layer and determined that there is still a problem, check the IS-IS layer, verify that IS-IS adjacencies are up, and make sure that the interfaces and IS-IS protocol are configured correctly.

Figure 32: MPLS Network Broken at the IS-IS Protocol Layer



To check the IS-IS protocol, follow these steps:

1. [Verify the LSP on page 534](#)
2. [Verify IS-IS Adjacencies and Interfaces on page 535](#)
3. [Verify the IS-IS Configuration on page 536](#)
4. [Take Appropriate Action on page 537](#)
5. [Verify the LSP Again on page 538](#)

Verify the LSP

Purpose Confirm that interfaces are configured for IS-IS, that the IS-IS protocol is configured correctly, and that adjacencies are established.

Action To verify the label-switched path (LSP), enter the following command on the ingress, transit, and egress routers:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn,  ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    24 Oct 21 13:48:01  No Route toward dest [3 times]
    23 Oct 21 13:47:44  Deselected as active
    22 Oct 21 13:47:43  No Route toward dest[2 times]
    21 Oct 21 13:47:43  ResvTear received
    20 Oct 21 13:47:43  Down
    19 Oct 21 13:47:43  10.1.13.2: No Route toward dest[2 times]
    18 Oct 21 13:47:38  Record Route:  10.1.13.2 10.1.36.2
    [...Output truncated...]
  Created: Tue Oct 19 21:22:53 2004
Total 1 displayed, Up 0,  Down1

Egress LSP: 0 sessions
Total 0 displayed, Up 0,  Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0,  Down 0
```

Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0,  Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0,  Down 0
```

```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 3

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn,  ActiveRoute: 0 , LSPname:  R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 3 second(s).
    13 Oct 21 14:23:33 CSPF failed: no route toward 10.0.0.1[90 times]
    12 Oct 21 13:39:56 Deselected as active
    11 Oct 21 13:39:56 CSPF: could not determine self
    [...Output truncated...]
  Created: Tue Oct 19 22:28:30 2004
Total 1 displayed, Up 0,  Down1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning The sample output shows that LSP **R1-to-R6** and the reverse LSP **R6-to-R1** are down, and there are no LSP sessions on transit router R3.

Verify IS-IS Adjacencies and Interfaces

Purpose When you check the IS-IS layer, you verify that IS-IS adjacencies are up and that the IS-IS interfaces are included at the protocol level.

Action To verify the functioning of adjacent interfaces, enter the following commands from the relevant routers:

```
user@host> show isis adjacency
user@host> show isis interface
```

Sample Output 1

```
user@R1> show isis adjacency
Interface          System      L State      HoId (secs) SNPA
so-0/0/0.0         R2          2 Up         20
so-0/0/1.0         R5          2 Up         23
so-0/0/2.0         R3          2 Up         26
```

```
user@R3> show isis adjacency
Interface          System      L State      HoId (secs) SNPA
so-0/0/0.0         R4          2 Up         23
so-0/0/1.0         R2          2 Up         21
so-0/0/2.0         R1          2 Up         19
so-0/0/3.0         R6          2 Down       0
```

```
user@R6> show isis adjacency
IS-IS instance is not running
```

Sample Output 2

```
user@R1> show isis interface
IS-IS interface database:
Interface          L  CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0              0   0x1  Passive              Passive          0/0
so-0/0/0.0         2   0x1  Disabled             Point to Point   10/10
so-0/0/1.0         2   0x1  Disabled             Point to Point   10/10
so-0/0/2.0         2   0x1  Disabled             Point to Point   10/10
```

```
user@R3> show isis interface
IS-IS interface database:
Interface          L  CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0              0   0x1  Passive              Passive          0/0
so-0/0/0.0         2   0x1  Disabled             Point to Point   10/10
so-0/0/1.0         2   0x1  Disabled             Point to Point   10/10
so-0/0/2.0         2   0x1  Disabled             Point to Point   10/10
so-0/0/3.0         2   0x1  Disabled             Point to Point   10/10
```

```
user@R6> show isis interface
IS-IS interface database:
Interface          L  CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0              0   0x1  Passive              Passive          0/0
so-0/0/0.0         1   0x1  Point to Point       Disabled         10/10
so-0/0/1.0         1   0x1  Down                 Disabled         10/10
so-0/0/2.0         1   0x1  Point to Point       Disabled         10/10
so-0/0/3.0         1   0x1  Point to Point       Disabled         10/10
```

Meaning Sample Output 1 shows that ingress router R1 has established adjacencies with the relevant routers. Transit router R3 does not have an adjacency with egress router R6, and egress router R6 has no adjacencies established in the network shown in *MPLS Network Broken at the IP and IGP Layers*, indicating that the problem might be at the IS-IS protocol level.

Sample Output 2 shows that R1 and R2 are Level 2 routers, in contrast to R6 which is a Level 1 router. When a router is configured explicitly as a Level 1 or Level 2 router, it does not communicate with routers configured at a different level. Level 1 routers communicate with other Level 1 routers within their area, while Level 2 routers communicate with other Level 2 routers, and toward other autonomous systems. Because all the routers in this network are configured for Level 2, they cannot form an adjacency with R6, which is incorrectly configured as a Level 1 router.

Verify the IS-IS Configuration

Purpose When you have determined that the problem is probably at the IS-IS protocol level, check the IS-IS configuration of the routers in your network.

Action To verify the IS-IS configuration, enter the following command from the relevant routers:

```
user@host> show configuration protocols isis
```

Sample Output

```
user@R1> show configuration protocols isis
level 1 disable;
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface lo0.0; {
    passive
```

```
user@R3> show configuration protocols isis
level 1 disable;
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive
```

```
user@R6> show configuration protocols isis
level 2 disable; <<< Incorrect level disabled
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive
```

Meaning The sample output shows that R6 has Level 2 disabled, while R1 and R3 have Level 1 disabled. For IS-IS adjacencies to establish, routers need to be at the same level. Another common configuration error is to omit the loopback interface (lo0) from the configuration at the **[edit protocols isis]** hierarchy level. IS-IS does not function correctly if the loopback interface (lo0) is not configured at this level. In addition, including the **passive** statement ensures that protocols are not run over the loopback interface (lo0) and that the loopback interface (lo0) is advertised correctly throughout the network.

Take Appropriate Action

Problem Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In the example below, the routers are configured to function at different levels of the IS-IS protocol.

Solution To correct the error in this example, enter the following commands:

Sample Output

```
[edit protocols isis]
user@R6# show
level 2 disable;
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive

[edit protocols isis]
user@R6# delete level 2

[edit protocols isis]
user@R6# set level 1 disable

[edit protocols isis]
user@R6# show
level 1 disable;
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive

[edit protocols isis]
user@R6# commit
commit complete

[edit protocols isis]
user@R6# run show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R5          2 Up         22
so-0/0/1.0         R4          2 Up         22
so-0/0/2.0         R2          2 Up         22
so-0/0/3.0         R3          2 Up         22
```

Meaning The sample output shows that the configuration error on egress router R6 has been corrected, and IS-IS adjacencies are now established.

Verify the LSP Again

Purpose After taking the appropriate action to correct the error, the label-switched path (LSP) needs to be checked again to confirm that the problem in the RSVP layer has been resolved.

Action To verify that the LSP is up and traversing the network as expected, enter the following command from the ingress, egress, and transit routers:

```
user@host> show mpls lsp extensive
```


Sample Output 1

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up,  ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
  10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.13.2 10.1.36.2
      5 Oct 21 15:52:07 Selected as active path
      4 Oct 21 15:52:07 Record Route: 10.1.13.2 10.1.36.2
      3 Oct 21 15:52:07 Up
      2 Oct 21 15:52:07 Originate Call
      1 Oct 21 15:52:07 CSPF: computation result accepted
    Created: Thu Oct 21 15:52:06 2004
  Total 1 displayed,  Up1 , Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-to-R1 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 142, Since: Thu Oct 21 15:41:59 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39082 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 17 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
  Total 1 displayed,  Up1 , Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 2

```

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 2 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
  LSPname: R6-to-R1 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -

```

```

Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100528, Label out: 3
Time left: 125, Since: Thu Oct 21 15:29:26 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 39082 protocol 0
PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 17 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.13.1 (so-0/0/2.0) 17 pkts
RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 17 pkts
Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1

```

10.0.0.6

```

From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100544, Label out: 3
Time left: 147, Since: Thu Oct 21 15:39:33 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47963 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 4 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 4 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 4 pkts
Explct route: 10.1.36.2
Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed, Up 2, Down 0

```

Sample Output 3

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

```

10.0.0.1

```

From: 10.0.0.6, State: Up, ActiveRoute: 1,  LSPname: R6-to-R1
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1S 10.1.13.1S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.36.1 10.1.13.1
18 Oct 21 15:34:18 Selected as active path
17 Oct 21 15:34:17 Record Route: 10.1.36.1 10.1.13.1
16 Oct 21 15:34:17 Up
15 Oct 21 15:34:17 Originate Call
14 Oct 21 15:34:17 CSPF: computation result accepted
[...Output truncated...]
Created: Tue Oct 19 22:28:30 2004
Total 1 displayed, Up 1, Down 0

```

```
Egress LSP: 1 sessions
```

10.0.0.6

```

From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -

```

```

Resv style: 1 FF, Label in: 3, Label out: -
Time left: 126, Since: Thu Oct 21 15:44:25 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47963 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 4 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```


Meaning Sample Outputs 1 and 3 from ingress router R1 and egress router R6 show that the LSP is now traversing the network along the expected path, from R1 through R3 to R6, and the reverse LSP, from R6 through R3 to R1. In addition, Sample Output 2 from transit router R3 shows that there are two transit LSP sessions, one from R1 to R6, and the other from R6 to R1.

Verifying the IS-IS Configuration on a Router in a Network

Purpose For IS-IS to run on a router (intermediate system) in your network, you must enable IS-IS on the router, configure a network entity title (NET) on the loopback interface (lo0), and configure **family iso** on all interfaces on which you want to run IS-IS. When you enable IS-IS on a router, Level 1 and Level 2 are enabled by default.

Figure 33 on page 541 illustrates an example of routers at different levels in an IS-IS topology.

Figure 33: Levels in an IS-IS Network Topology




The network in [Figure 33 on page 541](#) is organized hierarchically and consists of Level 2, Level 1/Level 2, and Level 1 routers in one autonomous system (AS) divided into four areas: 49.0001, 49.0002, 49.0003, and 49.0004. The Level 2 routers route toward other autonomous systems. The Level 1/Level 2 routers route between areas and to other autonomous systems. The Level 1 routers route within an area, and when the destination is outside the local area, they route toward a Level 1/Level 2 system.

In the following topics, the configuration of the various types of routers is examined.

[Figure 34 on page 542](#) provides more details about the IS-IS network topology in [Figure 33 on page 541](#) so that you can verify the configuration output of the various routers.

Figure 34: IS-IS Network Topology with Details



To verify that IS-IS is configured correctly on routers at different levels, follow these steps:

1. [Check the Configuration of a Level 1/Level 2 Router on page 542](#)
2. [Check the Configuration of a Level 1 Router on page 544](#)
3. [Check the Configuration of a Level 2 Router on page 545](#)

Check the Configuration of a Level 1/Level 2 Router

Purpose Check the configuration of a Level 1/Level 2 router.

Action To verify the IS-IS configuration of a Level 1/Level 2 router in your network, enter the following Junos OS command-line interface (CLI) commands:

```

user@host# [edit protocols isis] show
user@host# [edit protocols isis]
user@host# run show isis interface
user@host# [edit] edit interfaces
user@host# [edit interfaces] show

```

The following output is for an IS-IS configuration on R2, a Level 1/Level 2 router in the network shown.

Sample Output

```

[edit protocols isis]
user@R2# show
interface so-0/0/0.0 {
    level 2 metric 10;
    level 1 disable;
}
interface so-0/0/1.0 {
    level 2 disable;
    level 1 metric 10;
}
interface so-0/0/2.0 {
    level 2 metric 10;
    level 1 disable;
}
interface fxp0.0 {
    disable;
}
interface lo0.0;

[edit protocols isis]
user@R2# run show isis interface
IS-IS interface database:

```

Interface	L	CirID	Level 1 DR	Level 2 DR	L1/L2 Metric
lo0.0	0	0x1	Passive	Passive	0/0
so-0/0/0.0	2	0x1	Disabled	Point to Point	10/10
so-0/0/1.0	3	0x1	Point to Point	Point to Point	10/10
so-0/0/2.0	2	0x1	Disabled	Point to Point	10/10

```

[edit interfaces]
user@R2# show
so-0/0/0 {
    unit 0 {
        family inet {
            address 10.1.12.2/30;
        }
        family iso;
    }
}
so-0/0/1 {
    unit 0 {
        family inet {
            address 10.1.23.1/30;
        }
        family iso;
    }
}
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.26.1/30;
        }
    }
}

```

```

        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.2/32;
        }
        family iso {
            address 49.0002.1000.0000.0002.00;
        }
    }
}

```

Meaning The sample output shows a basic configuration of IS-IS on R2, a Level 1/Level 2 router. The basic configuration is at the **[edit protocols isis]** and **[edit interfaces]** hierarchy levels.

At the **[edit protocols isis]** level, five interfaces are included: so-0/0/0, so-0/0/1, so-0/0/2, fxp0, and the loopback interface (lo0). Two interfaces, so-0/0/0.0 and so-0/0/2.0, have Level 1 disabled, making them Level 2 interfaces. One interface, so-0/0/1.0, has Level 2 disabled, making it a Level 1 interface. The management interface (fxp0) is disabled so that IS-IS packets are not sent over it, and the loopback interface (lo0) is included because it becomes a point of connection from the router to the IS-IS network.

At the **[edit interfaces]** hierarchy level, all of the interfaces included in the **[edit protocols isis]** hierarchy level are configured with **family iso**, and the loopback interface (lo0) is configured with the NET address 49.0002.1000.0000.0002.00. Every router in an IS-IS network must have at least one NET address that identifies a point of connection to the IS-IS network. The NET address is generally configured on the loopback interface (lo0). Routers that participate in multiple areas can have multiple NET addresses.

Check the Configuration of a Level 1 Router

Purpose To check the configuration of a Level 1 router.

Action To check the configuration of a Level 1 router, enter the following CLI commands:

```

user@host# [edit protocols isis] show
user@host# [edit protocols isis] run show isis interface
user@host# [edit] edit interfaces
user@host# [edit interfaces] show

```

The following sample output is for R4, a Level 1 router in the network shown in The following output is for an IS-IS configuration on R2, a Level 1/Level 2 router in the network shown.

Sample Output

```

[edit protocols isis]

user@R4# show
level 2 disable;
interface so-0/0/2.0 {
    level 1 metric 10;
}

```

```

interface fxp0.0 {
    disable;
}
interface lo0.0;
[edit protocols isis]

user@R4# run show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0              0   0x1 Passive             Passive          0/0
so-0/0/2.0         1   0x1 Point to Point     Disabled         10/10
[edit interfaces]
user@R4# show
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.45.1/30;
        }
        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.4/32;
        }
        family iso {
            address 49.0001.1000.0000.0004.00;
        }
    }
}

```

Meaning The sample output shows a basic configuration of IS-IS on R4, a Level 1 router. The basic configuration is at the **[edit protocols isis]** and **[edit interfaces]** hierarchy levels.

At the **[edit protocols isis]** hierarchy level, three interfaces are included: so-0/0/2.0, fxp0, and the loopback interface (lo0). Level 2 is disabled on the router, making it a Level 1 router that sends packets within its local area, 49.0001. When a packet destination is outside the local area, R4 establishes an adjacency with the nearest Level 1/Level 2 router (R5) that forwards the packets. For more information about adjacencies, see [“Displaying the Status of IS-IS Adjacencies” on page 547](#).

One interface, so-0/0/2.0, is configured for IS-IS. The management interface (fxp0) is disabled so that IS-IS packets are not sent over it, and the loopback interface (lo0) is included because it becomes a point of connection from the router to the IS-IS network.

At the **[edit interfaces]** hierarchy level, the interface included in the **[edit protocols isis]** hierarchy level is also configured with **family iso**, and the loopback interface (lo0) is configured with the NET address of 49.0001.1000.0000.0004.00. Every router in an IS-IS network must have at least one NET address that identifies a point of connection to the IS-IS network. The NET address is generally configured on the loopback interface (lo0). Routers that participate in multiple areas can have multiple NET addresses.

Check the Configuration of a Level 2 Router

Purpose Check the configuration of a Level 2 router.

Action To check the configuration of a Level 2 router, enter the following CLI commands:

```
user@host# [edit protocols isis] show
user@host# [edit protocols isis] run show isis interface
user@host# [edit] edit interfaces
user@host# [edit interfaces] show
```

The following sample output is for R6, a Level 2 router in the network shown.

Sample Output

```
[edit protocols isis]
user@R6# show
level 1 disable;
interface so-0/0/0.0 {
    level 2 metric 10;
}
interface so-0/0/2.0 {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0;
```

```
[edit protocols isis]
user@R6# run show isis interface
IS-IS interface database:
Interface          L CirID Level1 DR          Level 2 DR          L1/L2 Metric
lo0.0              0  0x1 Passive              Passive              0/0
so-0/0/0.0         2  0x1 Disabled            Point to Point      10/10
so-0/0/2.0         2  0x1 Disabled            Point to Point      10/10
```

```
[edit interfaces]
user@R6# show
so-0/0/0 {
    unit 0 {
        family inet {
            address 10.1.56.2/30;
        }
        family iso;
    }
}
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.26.2/30;
        }
        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.6/32;
        }
        family iso {
            address 49.0003.1000.0000.0006.00;
        }
    }
}
```


Meaning The sample output shows a basic configuration of IS-IS on R6, a Level 2 router. The basic configuration is at the **[edit protocols isis]** and **[edit interfaces]** hierarchy levels.

At the **[edit protocols isis]** level, four interfaces are included: so-0/0/0.0, so-0/0/2.0, fxp0, and the loopback interface (lo0). Level 1 is disabled on the two SONET/SDH interfaces, making this a Level 2 router that routes between areas and toward other ASs. The management interface (fxp0) is disabled so that IS-IS packets are not sent over it, and the loopback interface (lo0) is included because it becomes a point of connection from the router to the IS-IS network.

At the **[edit interfaces]** hierarchy level, the interfaces included in the **[edit protocols isis]** hierarchy level are also configured with **family iso**, and the loopback interface (lo0) is configured with the NET address of 49.0003.1000.0000.0006.00. Every router in an IS-IS network must have at least one NET address that identifies a point of connection to the IS-IS network. The NET address is generally configured on the loopback interface (lo0). Routers that participate in multiple areas can have multiple NET addresses.

Related Documentation


- [Example: Configuring Multi-Level IS-IS on page 19](#)

Displaying the Status of IS-IS Adjacencies

Purpose Assuming that all the routers are correctly configured for IS-IS, you can verify which neighbors are adjacent and able to exchange IS-IS data. In addition, you can examine the set of routes installed in the forwarding table to verify that the routing protocol process (rpd) has relayed the correct information into the forwarding table.

[Figure 35 on page 548](#) illustrates the example IS-IS topology used for the procedures in this topic.

Figure 35: IS-IS Network Topology



The network consists of Level 1 and Level 2 adjacencies. Level 1 adjacencies are within areas 49.0001 and 49.0002. Level 2 adjacencies occur between all directly connected Level 2 routers regardless of which area they are in. For example, R5 is in area 49.0001, R6 is in area 49.0003, R1 is in area 49.0004, and R2 is in area 49.0002. The network in [Figure 35 on page 548](#) should have the following adjacencies:

- Level 2 adjacencies between all directly connected Level 2 routers (R1, R2, R5, and R6).
- Level 1 adjacencies between routers in area 49.0001 (R4 and R5) and between routers in area 49.0002 (R2 and R3).

To verify that routers are adjacent and able to exchange IS-IS data, follow these steps:

Displaying Detailed IS-IS Protocol Information

Action To trace IS-IS messages in detail, follow these steps:

1. Configure the flag to display detailed IS-IS protocol messages.

```
[edit protocols isis traceoptions]
user@host# set flag hello detail
```

2. Verify the configuration.

```
user@host# show
```

For example:

```
[edit protocols isis traceoptions]
user@host# show
```

```
file isislog size 10k files 10;
flag hello detail;
```

3. Commit the configuration.

```
user@host# commit
```

4. View the contents of the file containing the detailed messages.

```
user@host# run show log filename
```

For example:

```
user@host# run show log isislog
```

```
Nov 29 23:17:50 trace_on: Tracing to "/var/log/isislog" started
Nov 29 23:17:50 Sending PTP IIH on so-1/1/1.0
Nov 29 23:17:53 Sending PTP IIH on so-1/1/0.0
Nov 29 23:17:54 Received PTP IIH, source id abc-core-01 on so-1/1/0.0
Nov 29 23:17:54   from interface index 11
Nov 29 23:17:54   max area 0, circuit type 12, packet length 4469
Nov 29 23:17:54   hold time 30, circuit id 6
Nov 29 23:17:54   neighbor state up
Nov 29 23:17:54   speaks IP
Nov 29 23:17:54   area address 99.0008 (1)
Nov 29 23:17:54   IP address 10.10.10.29
Nov 29 23:17:54   4396 bytes of total padding
Nov 29 23:17:54   updating neighbor abc-core-01
Nov 29 23:17:55 Received PTP IIH, source id abc-core-02 on so-1/1/1.0
Nov 29 23:17:55   from interface index 12
Nov 29 23:17:55   max area 0, circuit type 12, packet length 4469
Nov 29 23:17:55   hold time 30, circuit id 6
Nov 29 23:17:55   neighbor state up
Nov 29 23:17:55   speaks IP
Nov 29 23:17:55   area address 99.0000 (1)
Nov 29 23:17:55   IP address 10.10.10.33
Nov 29 23:17:55   4396 bytes of total padding
Nov 29 23:17:55   updating neighbor abc-core-02
```

Meaning [Table 38 on page 549](#) lists tracing flags that can be configured specific to IS-IS and presents example output for some of the flags.

Table 38: IS-IS Protocol Tracing Flags

Tracing Flags	Description	Example Output
csn	Complete sequence number PDU (CSNP)	<p>Nov 28 20:02:48 Sending L2 CSN on interface so-1/1/0.0Nov 28 20:02:48 Sending L2 CSN on interface so-1/1/1.0</p> <p>With the detail option.</p> <p>Nov 28 20:06:08 Sending L2 CSN on interface so-1/1/1.0Nov 28 20:06:08 LSP abc-core-01.00-00 lifetime 1146Nov 28 20:06:08 sequence 0x1c4f8 checksum 0xa1e9Nov 28 20:06:08 LSP abc-core-02.00-00 lifetime 411Nov 28 20:06:08 sequence 0x7435 checksum 0x5424Nov 28 20:06:08 LSP abc-brdr-01.00-00 lifetime 465Nov 28 20:06:08 sequence 0xf73 checksum 0xab10Nov 28 20:06:08 LSP abc-edge-01.00-00 lifetime 1089Nov 28 20:06:08 sequence 0x1616 checksum 0xdb29Nov 28 20:06:08 LSP abc-edge-02.00-00 lifetime 1103Nov 28 20:06:08 sequence 0x45cc checksum 0x6883</p>

Table 38: IS-IS Protocol Tracing Flags (*continued*)

Tracing Flags	Description	Example Output
hello	Hello packet	Nov 28 20:13:50 Sending PTP IIH on so-1/1/1.0Nov 28 20:13:50 Received PTP IIH, source id abc-core-01 on so-1/1/0.0Nov 28 20:13:53 Received PTP IIH, source id abc-core-02 on so-1/1/1.0Nov 28 20:13:57 Sending PTP IIH on so-1/1/0.0Nov 28 20:13:58 Received PTP IIH, source id abc-core-01 on so-1/1/0.0Nov 28 20:13:59 Sending PTP IIH on so-1/1/1.0
lsp	Link-state PDUs (LSPs)	Nov 28 20:15:46 Received L2 LSP abc-edge-01.00-00, interface so-1/1/0.0Nov 28 20:15:46 from abc-core-01Nov 28 20:15:46 sequence 0x1617, checksum 0xd92a, lifetime 1197Nov 28 20:15:46 Updating L2 LSP abc-edge-01.00-00 in TEDNov 28 20:15:47 Received L2 LSP abc-edge-01.00-00, interface so-1/1/1.0Nov 28 20:15:47 from abc-core-02Nov 28 20:15:47 sequence 0x1617, checksum 0xd92a, lifetime 1197
lsp-generation	Link-state PDU generation packets	Nov 28 20:21:24 Regenerating L1 LSP abc-edge-03.00-00, old sequence 0x682Nov 28 20:21:27 Rebuilding L1, fragment abc-edge-03.00-00Nov 28 20:21:27 Rebuilt L1 fragment abc-edge-03.00-00, size 59Nov 28 20:31:52 Regenerating L2 LSP abc-edge-03.00-00, old sequence 0x689Nov 28 20:31:54 Rebuilding L2, fragment abc-edge-03.00-00Nov 28 20:31:54 Rebuilt L2 fragment abc-edge-03.00-00, size 256Nov 28 20:34:05 Regenerating L1 LSP abc-edge-03.00-00, old sequence 0x683Nov 28 20:34:08 Rebuilding L1, fragment abc-edge-03.00-00Nov 28 20:34:08 Rebuilt L1 fragment abc-edge-03.00-00, size 59
packets	All IS-IS protocol packets	Not available.
psn	Partial sequence number PDU (PSNP) packets	Nov 28 20:40:39 Received L2 PSN, source abc-core-01, interface so-1/1/0.0Nov 28 20:40:39 Received L2 PSN, source abc-core-02, interface so-1/1/1.0Nov 28 20:41:36 Sending L2 PSN on interface so-1/1/1.0Nov 28 20:41:36 Sending L2 PSN on interface so-1/1/0.0Nov 28 20:42:35 Received L2 PSN, source abc-core-02, interface so-1/1/1.0Nov 28 20:42:35 LSP abc-edge-03.00-00 lifetime 1196Nov 28 20:42:35 sequence 0x68c checksum 0x746dNov 28 20:42:35 Received L2 PSN, source abc-core-01, interface so-1/1/0.0Nov 28 20:42:35 LSP abc-edge-03.00-00 lifetime 1196Nov 28 20:42:35 sequence 0x68c checksum 0x746dNov 28 20:42:49 Sending L2 PSN on interface so-1/1/1.0Nov 28 20:42:49 LSP abc-core-01.00-00 lifetime 1197Nov 28 20:42:49 sequence 0x1c4fb checksum 0x9becNov 28 20:42:49 Sending L2 PSN on interface so-1/1/0.0Nov 28 20:42:49 LSP abc-core-01.00-00 lifetime 1197Nov 28 20:42:49 sequence 0x1c4fb checksum 0x9bec
spf	Shortest-path-first (SPF) calculations	Nov 28 20:44:01 Scheduling SPF for L1: ReconfigNov 28 20:44:01 Scheduling multicast SPF for L1: ReconfigNov 28 20:44:01 Scheduling SPF for L2: ReconfigNov 28 20:44:01 Scheduling multicast SPF for L2: ReconfigNov 28 20:44:02 Running L1 SPFNov 28 20:44:02 L1 SPF initialization complete: 0.000099s cumulative timeNov 28 20:44:02 L1 SPF primary processing complete: 0.000303s cumulative timeNov 28 20:44:02 L1 SPF result postprocessing complete: 0.000497s cumulative timeNov 28 20:44:02 L1 SPF RIB postprocessing complete: 0.000626s cumulative timeNov 28 20:44:02 L1 SPF routing table postprocessing complete: 0.000736s cumulative time

Related Documentation • [Example: Configuring Multi-Level IS-IS on page 19](#)

Analyzing IS-IS Link-State PDUs in Detail

To analyze IS-IS link-state PDUs in detail, follow these steps:

1. Configure IS-IS open messages.

```
[edit protocols isis traceoptions]
user@host# set flag lsp detail
```

2. Verify the configuration.

```
user@host# show
```

For example:

```
[edit protocols isis traceoptions]
user@host# show
file isislog size 5m world-readable;
flag error;
flag lsp detail;
```

3. Commit the configuration.

```
user@host# commit
```

4. View the contents of the file containing the detailed messages.

```
user@host# run show log filename
```

For example:

```
user@host# run show log isislog
Nov 28 20:17:24 Received L2 LSP abc-core-01.00-00, interface so-1/1/0.0
Nov 28 20:17:24   from abc-core-01
Nov 28 20:17:24   sequence 0x1c4f9, checksum 0x9fea, lifetime 1199
Nov 28 20:17:24   max area 0, length 426
Nov 28 20:17:24   no partition repair, no database overload
Nov 28 20:17:24   IS type 3, metric type 0
Nov 28 20:17:24   area address 99.0908 (1)
Nov 28 20:17:24   speaks CLNP
Nov 28 20:17:24   speaks IP
Nov 28 20:17:24   dyn hostname abc-core-01
Nov 28 20:17:24   IP address 10.10.134.11
Nov 28 20:17:24   IP prefix: 10.10.10.0/30 metric 1 up
Nov 28 20:17:24   IP prefix: 10.10.10.4/30 metric 5 up
Nov 28 20:17:24   IP prefix: 10.10.10.56/30 metric 5 up
Nov 28 20:17:24   IP prefix: 10.10.10.52/30 metric 1 up
Nov 28 20:17:24   IP prefix: 10.10.10.64/30 metric 5 up
Nov 28 20:17:24   IP prefix: 10.10.10.20/30 metric 5 up
Nov 28 20:17:24   IP prefix: 10.10.10.28/30 metric 5 up
Nov 28 20:17:24   IP prefix: 10.10.10.44/30 metric 5 up
Nov 28 20:17:24   IP prefix 10.10.10.0 255.255.255.252
Nov 28 20:17:24   internal, metrics: default 1
Nov 28 20:17:24   IP prefix 10.10.10.4 255.255.255.252
Nov 28 20:17:24   internal, metrics: default 5
Nov 28 20:17:24   IP prefix 10.10.10.56 255.255.255.252
Nov 28 20:17:24   internal, metrics: default 5
Nov 28 20:17:24   IP prefix 10.10.10.52 255.255.255.252
Nov 28 20:17:24   internal, metrics: default 1
Nov 28 20:17:24   IP prefix 10.10.10.64 255.255.255.252
Nov 28 20:17:24   internal, metrics: default 5
Nov 28 20:17:24   IP prefix 10.10.10.20 255.255.255.252
Nov 28 20:17:24   internal, metrics: default 5
Nov 28 20:17:24   IP prefix 10.10.10.28 255.255.255.252
Nov 28 20:17:24   internal, metrics: default 5
Nov 28 20:17:24   IP prefix 10.10.10.44 255.255.255.252
Nov 28 20:17:24   internal, metrics: default 5
Nov 28 20:17:24   IS neighbors:
Nov 28 20:17:24   IS neighbor abc-core-02.00
```

```
Nov 28 20:17:24      internal, metrics: default 1
[...Output truncated...]
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IS neighbor abc-brdr-01.00
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IS neighbor abc-core-02.00, metric: 1
Nov 28 20:17:24      IS neighbor abc-esr-02.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-edge-03.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-edge-01.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-edge-02.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-brdr-01.00, metric: 5
Nov 28 20:17:24      IP prefix: 10.10.134.11/32 metric 0 up
Nov 28 20:17:24      IP prefix: 10.11.0.0/16 metric 5 up
Nov 28 20:17:24      IP prefix: 10.211.0.0/16 metric 0 up
Nov 28 20:17:24      IP prefix 10.10.134.11 255.255.255.255
Nov 28 20:17:24      internal, metrics: default 0
Nov 28 20:17:24      IP prefix 10.11.0.0 255.255.0.0
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IP prefix 10.211.0.0 255.255.0.0
Nov 28 20:17:24      internal, metrics: default 0
Nov 28 20:17:24      Updating LSP
Nov 28 20:17:24      Updating L2 LSP abc-core-01.00-00 in TED
Nov 28 20:17:24      Analyzing subtlv's for abc-core-02.00
Nov 28 20:17:24      Analysis complete
Nov 28 20:17:24      Analyzing subtlv's for abc-esr-02.00
Nov 28 20:17:24      Analysis complete
Nov 28 20:17:24      Analyzing subtlv's for abc-edge-03.00
Nov 28 20:17:24      Analysis complete
Nov 28 20:17:24      Analyzing subtlv's for abc-edge-01.00
Nov 28 20:17:24      Analysis complete
Nov 28 20:17:24      Analyzing subtlv's for abc-edge-02.00
Nov 28 20:17:24      Analysis complete
Nov 28 20:17:24      Analyzing subtlv's for abc-brdr-01.00
Nov 28 20:17:24      Analysis complete
Nov 28 20:17:24      Scheduling L2 LSP abc-core-01.00-00 sequence 0x1c4f9 on
interface so-1/1/1.0
```

Related Documentation • [Example: Configuring Multi-Level IS-IS on page 19](#)

Displaying Sent or Received IS-IS Protocol Packets

To configure the tracing for only sent or received IS-IS protocol packets, follow these steps:

1. Configure the flag to display sent, received, or both sent and received packets.

```
[edit protocols isis traceoptions]
user@host# set flag hello send
```

or

```
[edit protocols isis traceoptions]
user@host# set flag hello receive
```

or

```
[edit protocols isis traceoptions]
user@host# set flag hello
```

2. Verify the configuration.

```
user@host# show
```

For example:

```
[edit protocols isis traceoptions]
user@host# show
file isislog size 10k files 10;
flag hello send;
```

or

```
[edit protocols isis traceoptions]
user@host# show
file isislog size 10k files 10;
flag hello receive;
```

or

```
[edit protocols isis traceoptions]
user@host# show
file isislog size 10k files 10;
flag hello send receive;
```

3. Commit the configuration.

```
user@host# commit
```

4. View the contents of the file containing the detailed messages.

```
user@host# run show log filename
```

For example:

```
user@host# run show log isislog
Sep 27 18:17:01 ISIS periodic xmit to 01:80:c2:00:00:15 (IFL 2)
Sep 27 18:17:01 ISIS periodic xmit to 01:80:c2:00:00:14 (IFL 2)
Sep 27 18:17:03 ISIS periodic xmit to 01:80:c2:00:00:15 (IFL 2)
Sep 27 18:17:04 ISIS periodic xmit to 01:80:c2:00:00:14 (IFL 2)
Sep 27 18:17:06 ISIS L2 hello from 0000.0000.0008 (IFL 2) absorbed
Sep 27 18:17:06 ISIS periodic xmit to 01:80:c2:00:00:15 (IFL 2)
Sep 27 18:17:06 ISIS L1 hello from 0000.0000.0008 (IFL 2) absorbed
```

Related Documentation

- [Example: Configuring Multi-Level IS-IS on page 19](#)

Routing Protocol Process Memory FAQs

- [Routing Protocol Process Memory FAQs Overview on page 555](#)
- [Routing Protocol Process Memory FAQs on page 556](#)

Routing Protocol Process Memory FAQs Overview

Junos OS is based on the FreeBSD Unix operating system. The open source software is modified and hardened to operate in the device's specialized environment. For example, some executables have been deleted, while other utilities were de-emphasized. Additionally, certain software processes were added to enhance the routing functionality. The result of this transformation is the kernel, the heart of the Junos OS software.

The kernel is responsible for operating multiple processes that perform the actual functions of the device. Each process operates in its own protected memory space, while the communication among all the processes is still controlled by the kernel. This separation provides isolation between the processes, and resiliency in the event of a process failure. This is important in a core routing platform because a single process failure does not cause the entire device to cease functioning.

Some of the common software processes include the routing protocol process (rpd) that controls the device's protocols, the device control process (dcd) that controls the device's interfaces, the management process (mgd) that controls user access to the device, the chassis process (chassisd) that controls the device's properties itself, and the Packet Forwarding Engine process (pfed) that controls the communication between the device's Packet Forwarding Engine and the Routing Engine. The kernel also generates specialized processes as needed for additional functionality, such as SNMP, the Virtual Router Redundancy Protocol (VRRP), and Class of Service (CoS).

The routing protocol process is a software process within the Routing Engine software, which controls the routing protocols that run on the device. Its functionality includes all protocol messages, routing table updates, and implementation of routing policies.

The routing protocol process starts all configured routing protocols and handles all routing messages. It maintains one or more routing tables, which consolidate the routing information learned from all routing protocols. From this routing information, the routing protocol process determines the active routes to network destinations and installs these routes into the Routing Engine's forwarding table. Finally, it implements routing policy, which allows you to control the routing information that is transferred between the routing

protocols and the routing table. Using routing policy, you can filter and limit the transfer of information as well as set properties associated with specific routes.

Related Documentation

- [Routing Protocol Process Memory FAQs on page 556](#)

Routing Protocol Process Memory FAQs

The following sections present the most frequently asked questions and answers related to the routing protocol process memory utilization, operation, interpretation of related command outputs, and troubleshooting the software process.

Frequently Asked Questions: Routing Protocol Process Memory

This section presents frequently asked questions and answers related to the memory usage of the routing protocol process.

Why does the routing protocol process use excessive memory?

The routing protocol process uses hundreds of megabytes of RAM in the Routing Engine to store information needed for the operation of routing and related protocols, such as BGP, OSPF, IS-IS, RSVP, LDP and MPLS. Such huge consumption of memory is common for the process, as the information it stores includes routes, next hops, interfaces, routing policies, labels, and label-switched paths (LSPs). Because access to the RAM memory is much faster than access to the hard disk, most of the routing protocol process information is stored in the RAM memory instead of using the hard disk space. This ensures that the performance of the routing protocol process is maximized.

How can I check the amount of memory the routing protocol process is using?

You can check routing protocol process memory usage by entering the **show system processes** and the **show task memory** Junos OS command-line interface (CLI) operational mode commands.

The **show system processes** command displays information about software processes that are running on the device and that have controlling terminals. The **show task memory** command displays memory utilization for routing protocol tasks on the Routing Engine.

You can check the routing protocol process memory usage by using the **show system processes** command with the **extensive** option. The **show task memory** command displays a report generated by the routing protocol process on its own memory usage. However, this report does not display all the memory used by the process. The value reported by the routing protocol process does not account for the memory used for the **TEXT** and **STACK** segments, or the memory used by the process's internal memory manager. Further, the Resident Set Size value includes shared library pages used by the routing protocol process.

For more information about checking the routing protocol process memory usage.

For more information, see the **show system processes** command and the **show task memory** command.

I just deleted a large number of routes from the routing protocol process. Why is it still using so much memory?

The **show system processes extensive** command displays a **RES** value measured in kilobytes. This value represents the amount of program memory resident in the physical memory. This is also known as RSS or Resident Set Size. The **RES** value includes shared library pages used by the process. Any amount of memory freed by the process might still be considered part of the **RES** value. Generally, the kernel delays the migrating of memory out of the **Inact** queue into the **Cache** or **Free** list unless there is a memory shortage. This can lead to large discrepancies between the values reported by the routing protocol process and the kernel, even after the routing protocol process has freed a large amount of memory.

Frequently Asked Questions: Interpreting Routing Protocol Process-Related Command Outputs

This section presents frequently asked questions and answers about the routing protocol process-related Junos OS command-line interface (CLI) command outputs that are used to display the memory usage of the routing protocol process.

How do I interpret memory numbers displayed in the show system processes extensive command output?

The **show system processes extensive** command displays exhaustive system process information about software processes that are running on the device and have controlling terminals. This command is equivalent to the UNIX **top** command. However, the UNIX **top** command shows real-time memory usage, with the memory values constantly changing, while the **show system processes extensive** command provides a snapshot of memory usage in a given moment.

To check overall CPU and memory usage, enter the **show system processes extensive** command. Refer to [Table 39 on page 558](#) for information about the **show system processes extensive** commands output fields.

```
user@host> show system processes extensive
last pid: 544; load averages: 0.00, 0.00, 0.00 18:30:33
37 processes: 1 running, 36 sleeping

Mem: 25M Active, 3968K Inact, 19M Wired, 184K Cache, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

  PID USERNAME PRI NICE SIZE RES STATE TIME WCPU CPU COMMAND
    544 root    30  0  604K 768K RUN   0:00 0.00% 0.00% top
      3 root    28  0    0K  12K psleep 0:00 0.00% 0.00% vmdaemon
      4 root    28  0    0K  12K update 0:03 0.00% 0.00% update
    528 aviva    18  0  660K 948K pause 0:00 0.00% 0.00% tcsh
    204 root    18  0  300K 544K pause 0:00 0.00% 0.00% csh
    131 root    18  0  332K 532K pause 0:00 0.00% 0.00% cron
    186 root    18  0  196K  68K pause 0:00 0.00% 0.00% watchdog
     27 root    10  0  512M 16288K mfsidl 0:00 0.00% 0.00% mount_mfs
      1 root    10  0  620K 344K wait 0:00 0.00% 0.00% init
    304 root     3  0  884K 900K ttyin 0:00 0.00% 0.00% bash
    200 root     3  0  180K 540K ttyin 0:00 0.00% 0.00% getty
    203 root     3  0  180K 540K ttyin 0:00 0.00% 0.00% getty
    202 root     3  0  180K 540K ttyin 0:00 0.00% 0.00% getty
    201 root     3  0  180K 540K ttyin 0:00 0.00% 0.00% getty
    194 root     2  0 2248K 1640K select 0:11 0.00% 0.00% rpd
    205 root     2  0   964K 800K select 0:12 0.00% 0.00% tnp.chassisd
```

```

189 root      2 -12  352K  740K select  0:03  0.00%  0.00% xntpd
114 root      2  0   296K  612K select  0:00  0.00%  0.00% amd
188 root      2  0   780K  600K select  0:00  0.00%  0.00% dcd
527 root      2  0   176K  580K select  0:00  0.00%  0.00% rlogind
195 root      2  0   212K  552K select  0:00  0.00%  0.00% inetd
187 root      2  0   192K  532K select  0:00  0.00%  0.00% tnetd
 83 root      2  0   188K  520K select  0:00  0.00%  0.00% syslogd
538 root      2  0  1324K  516K select  0:00  0.00%  0.00% mgd
 99 daemon    2  0   176K  492K select  0:00  0.00%  0.00% portmap
163 root      2  0   572K  420K select  0:00  0.00%  0.00% nsrexecd
192 root      2  0   560K  400K select  0:10  0.00%  0.00% snmpd
191 root      2  0  1284K  376K select  0:00  0.00%  0.00% mgd
537 aviva     2  0   636K  364K select  0:00  0.00%  0.00% cli
193 root      2  0   312K  204K select  0:07  0.00%  0.00% mib2d
  5 root      2  0      0K   12K pfesel  0:00  0.00%  0.00% if_pfe
  2 root     -18  0      0K   12K psleep  0:00  0.00%  0.00% pagedaemon
  0 root     -18  0      0K    0K sched   0:00  0.00%  0.00% swapper

```

Table 39 on page 558 describes the output fields that represent the memory values for the **show system processes extensive** command. Output fields are listed in the approximate order in which they appear.

Table 39: show system processes extensive Output Fields

Field Name	Field Description
Mem	Information about physical and virtual memory allocation.
Active	Memory allocated and actively used by the program.
Inact	Memory allocated but not recently used or memory freed by the programs. Inactive memory remains mapped in the address space of one or more processes and, therefore, counts toward the RSS value of those processes.
Wired	Memory that is not eligible to be swapped, usually used for in-kernel memory structures and/or memory physically locked by a process.
Cache	Memory that is not associated with any program and does not need to be swapped before being reused.
Buf	Size of memory buffer used to hold data recently called from the disk.
Free	Memory that is not associated with any programs. Memory freed by a process can become Inactive , Cache , or Free , depending on the method used by the process to free the memory.
Swap	Information about swap memory. <ul style="list-style-type: none"> • Total—Total memory available to be swapped to disk. • Used—Memory swapped to disk. • Free—Memory available for further swap.

The rest of the command output displays information about the memory usage of each process. The **SIZE** field indicates the size of the virtual address space, and the **RES** field indicates the amount of the program in physical memory, which is also known as RSS or Resident Set Size. For more information, see the **show system processes** command.

What is the difference between Active and Inact memory that is displayed by the show system processes extensive command?

When the system is under memory pressure, the pageout process reuses memory from the free, cache, inactive and, if necessary, active pages. When the pageout process runs, it scans memory to see which pages are good candidates to be unmapped and freed up. Thus, the distinction between **Active** and **Inact** memory is only used by the pageout process to determine which pool of pages to free first at the time of a memory shortage.

The pageout process first scans the **Inact** list, and checks whether the pages on this list have been accessed since the time they have been listed here. The pages that have been accessed are moved from the **Inact** list to the **Active** list. On the other hand, pages that have not been accessed become prime candidates to be freed by the pageout process. If the pageout process cannot produce enough free pages from the **Inact** list, pages from the **Active** list get freed up.

Because the pageout process runs only when the system is under memory pressure, the pages on the **Inact** list remain untouched – even if they have not been accessed recently – when the amount of **Free** memory is adequate.

How do I interpret memory numbers displayed in the show task memory command output?

The **show task memory** command provides a comprehensive picture of the memory utilization for routing protocol tasks on the Routing Engine. The routing protocol process is the main task that uses Routing Engine memory.

To check routing process memory usage, enter the **show task memory** command. Refer to [Table 40 on page 559](#) for information about the **show task memory** command output fields.

```
user@host> show task memory
Memory      Size (kB)  %Available  When
Currently In Use:    29417      3%         now
Maximum Ever Used:   33882      4%         00/02/11 22:07:03
Available:          756281    100%        now
```

[Table 40 on page 559](#) describes the output fields for the **show task memory** command. Output fields are listed in the approximate order in which they appear.

Table 40: show task memory Output Fields

Field Name	Field Description
Memory Currently In Use	Memory currently in use. Dynamically allocated memory plus the DATA segment memory in kilobytes.
Memory Maximum Ever Used	Maximum memory ever used.
Memory Available	Memory currently available.

The **show task memory** command does not display all the memory used by the routing protocol process. This value does not account for the memory used for the **TEXT** and

STACK segments, or the memory used by the routing protocol process's internal memory manager.

Why is the Currently In Use value less than the RES value?

The **show task memory** command displays a **Currently In Use** value measured in kilobytes. This value represents the memory currently in use. It is the dynamically allocated memory plus the **DATA** segment memory. The **show system processes extensive** command displays a **RES** value measured in kilobytes. This value represents the amount of program memory resident in the physical memory. This is also known as RSS or Resident Set Size.

The **Currently In Use** value does not account for all of the memory that the routing protocol process uses. This value does not include the memory used for the **TEXT** and the **STACK** segments, and a small percentage of memory used by the routing protocol process's internal memory manager. Further, the **RES** value includes shared library pages used by the routing protocol process.

Any amount of memory freed by the routing protocol process might still be considered part of the **RES** value. Generally, the kernel delays the migrating of memory out of the **Inact** queue into the **Cache** or **Free** list unless there is a memory shortage. This can lead to large discrepancies between the **Currently In Use** value and the **RES** value.

Frequently Asked Questions: Routing Protocol Process Memory Swapping

This section presents frequently asked questions and answers related to the memory swapping of the routing protocol process from the Routing Engine memory to the hard disk memory.

How do I monitor swap activity?

When the system is under memory pressure, the pageout process reuses memory from the free, cache, inact and, if necessary, active pages. You can monitor the swap activity by viewing the syslog message reported by the kernel during periods of high pageout activity.

The syslog message appears as follows:

```
Mar  3 20:08:02 olympic /kernel: High pageout rate!! 277 pages/sec.
```

You can use the **vmstat -s** command to print the statistics for the swapout activity. The displayed statistics appear as follows:

```
0 swap pager pageouts
0 swap pager pages paged out
```

The **swap pager pageouts** is the number of pageout operations to the swap device, and the **swap pager pages paged out** is the number of pages paged out to the swap device.

Why does the system start swapping when I try to dump core using the request system core-dumps command?

The **request system core-dumps** command displays a list of system core files created when the device has failed. This command can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification

date, path, and filename. You can use the **core-filename** option and the **core-file-info**, **brief**, and **detail** options to display more information about the specified core-dump files.

You can use the **request system core-dumps** command to perform a non-fatal core-dump without aborting the routing protocol process. To do this, the routing protocol process is forked, generating a second copy, and then aborted. This process can double the memory consumed by the two copies of the routing protocol processes, pushing the system into swap.

Why does the show system processes extensive command show that memory is swapped to disk although there is plenty of free memory?

Memory can remain swapped out indefinitely if it is not accessed again. Therefore, the **show system processes extensive** command shows that memory is swapped to disk even though there is plenty of free memory, and such a situation is not unusual.

Frequently Asked Questions: Troubleshooting the Routing Protocol Process

This section presents frequently asked questions and answers related to a shortage of memory and memory leakage by the routing protocol process.

What does the RPD_OS_MEMHIGH message mean?

The **RPD_OS_MEMHIGH** message is written into the system message file if the routing protocol process is running out of memory. This message alerts you that the routing protocol process is using the indicated amount and percentage of Routing Engine memory, which is considered excessive. This message is generated either because the routing protocol process is leaking memory or the use of system resources is excessive, perhaps because routing filters are misconfigured or the configured network topology is very complex.

When the memory utilization for the routing protocol process is using all available Routing Engine DRAM memory (Routing Engines with maximum 2 GB DRAM) or reaches the limit of 2 GB of memory (Routing Engines with 4 GB DRAM), a message of the following form is written every minute in the syslog message file:

RPD_OS_MEMHIGH: Using 188830 KB of memory, 100 percent of available

This message includes the amount, in kilobytes and/or the percentage, of the available memory in use.

This message should not appear under normal conditions, as any further memory allocations usually require a portion of existing memory to be written to swap. As a recommended solution, increase the amount of RAM in the Routing Engine. For more information, go to <http://kb.juniper.net/InfoCenter/index?page=content&id=KB14186>.

What can I do when there is a memory shortage even after a swap?

It is not recommended for the system to operate in this state, notwithstanding the existence of swap. The protocols that run in the routing protocol process usually have a real-time requirement that cannot reliably withstand the latency of being swapped to hard disk. If the memory shortage has not resulted from a memory leak, then either a

reduction in the memory usage or an upgrade to a higher memory-capacity Routing Engine is required.

How do I determine whether there is a memory leak in the routing protocol process?

Memory leaks are typically the result of a seemingly unbounded growth in the memory usage of a process as reported by the **show system processes extensive** command.

There are two classes of memory leaks that the routing protocol process can experience.

- The first class occurs when the allocated memory that is no longer in use is not freed. This class of leak can usually be fixed by taking several samples of the **show task memory detail** command over a period of time and comparing the deltas.
- The second class occurs when there is a late access to freed memory. If the access is not outside the mapped address space, the kernel backfills the accessed page with real memory. This backfill is done without the knowledge of the routing protocol process's internal memory allocator, which makes this class of leak much more difficult to resolve. If a memory leak of this class is suspected, writing the state of the system to a disk file (creating a core file) is suggested.

A large discrepancy between the **RES** value and the **Currently In Use** value might indicate a memory leak. However, large discrepancies can also occur for legitimate reasons. For example, the memory used for the **TEXT** and **STACK** segments or the memory used by the routing protocol process's internal memory manager might not be displayed. Further, the **RES** value includes shared library pages used by the process.

What is the task_timer?

The source of a routing protocol process memory leak can usually be identified by dumping the timers for each task. You can use the **show task task-name** command to display routing protocol tasks on the Routing Engine. Tasks can be baseline tasks performed regardless of the device's configuration, and other tasks that depend on the device configuration.

For more information, see the **show task** command.

Related Documentation

- [Routing Protocol Process Memory FAQs Overview on page 555](#)

PART 5

Index

- [Index on page 565](#)

Index

Symbols

#, comments in configuration statements.....	xix
(), in syntax descriptions.....	xix
< >, in syntax descriptions.....	xix
[], in configuration statements.....	xix
{ }, in configuration statements.....	xix
(pipe), in syntax descriptions.....	xix

A

addresses	
IP addresses.....	41
adjacencies	
IS-IS, verifying	535
advertisements, displaying	398
alias option for static-host-mapping	
statement.....	290
alias statement.....	290
authentication	
algorithm	
IS-IS.....	45
IS-IS.....	45
keychains	
IS-IS.....	45
authentication configuration	
BFD.....	101
authentication-algorithm statement	
IS-IS	
usage guidelines.....	46
authentication-key statement	
IS-IS.....	216
usage guidelines.....	43, 46
authentication-key-chain statement.....	217
IS-IS	
usage guidelines.....	46
authentication-type statement	
IS-IS.....	218
usage guidelines.....	43

B

backup-spf-options statement	
IS-IS.....	219

best routes, displaying.....	405
BFD	
authentication configuration.....	101
protocol.....	91
sessions	
clearing.....	306
displaying.....	327
BFD authentication	
for IS-IS.....	104
bfd-liveness-detection statement	
IS-IS.....	220
usage guidelines.....	91
BGP	
communities	
policy, routing.....	72
BGP (Border Gateway Protocol)	
export policy for CLNS.....	207
Bidirectional Forwarding Detection See BFD	
braces, in configuration statements.....	xix
brackets	
angle, in syntax descriptions.....	xix
square, in configuration statements.....	xix

C

checksum	
for IS-IS.....	51
checksum statement.....	222
clear bfd adaptation command.....	305
clear bfd session command.....	306
clear isis adjacency command.....	307
clear isis database command.....	309
clear isis overload command.....	311
clear isis statistics command.....	313
CLNP (Connectionless Network Protocol).....	3
CLNS (Connectionless Network Service) VPNs	
BGP export policy.....	207
IS-IS.....	207
CLNS node reachability, checking.....	315
clns-routing statement	
IS-IS.....	222
clns-updown-compatibility statement	
IS-IS.....	223
comments, in configuration statements.....	xix
communities	
policy, routing.....	72
community	
example.....	73
complete sequence number PDUs, IS-IS.....	6

conditions	
routing policy.....	386
connections	
testing	
CLNS connections.....	315
context-identifier	
IS-IS.....	224
conventions	
text and syntax.....	xviii
credibility-protocol-preference	
traffic engineering	
IS-IS.....	296
csn (tracing flag).....	293
csn, IS-IS protocol tracing flag.....	549
csnp-interval statement.....	225
usage guidelines.....	195
curly braces, in configuration statements.....	xix
customer support.....	xx
contacting JTAC.....	xx

D

default route	
configuring on logical systems.....	82
delay statement	
IS-IS.....	288
designated router, IS-IS	
about.....	39
configuring election priority.....	40
detail statement	
IS-IS protocol.....	548, 551
detection-time statement	
IS-IS.....	220
disable statement	
IS-IS.....	226
graceful restart.....	231
LDP synchronization.....	227
documentation	
comments on.....	xix
dual stacking	
IS-IS.....	126
dynamic overload bit, resetting for IS-IS.....	311

E

error (tracing flag)	
IS-IS.....	293
export route information, displaying.....	429
export statement	
IS-IS.....	228
usage guidelines.....	73

external-preference statement	
IS-IS.....	229

F

family statement	
IS-IS.....	230
usage guidelines.....	159
FAQs	
routing protocol process memory.....	555, 556
flag statements	
IS-IS protocol.....	548, 551, 552
font conventions.....	xviii
forwarding adjacencies	
IS-IS.....	173
forwarding table	
route entries, displaying.....	449

G

graceful-restart (tracing flag)	
IS-IS.....	293
graceful-restart statement	
IS-IS.....	231

H

hello	
IS-IS protocol tracing flag.....	550
statement	
IS-IS protocol.....	548, 552
hello (tracing flag)	
IS-IS.....	293
hello packets	
IS-IS.....	6
hello-authentication-key statement.....	232
hello-authentication-key-chain statement.....	233
hello-authentication-type statement.....	234
hello-interval statement	
IS-IS.....	235
hello-padding statement.....	236
helper-disable statement	
IS-IS.....	231
hidden routes, displaying.....	463
hold-time statement	
IS-IS.....	238
LDP synchronization.....	239
holddown statement	
IS-IS.....	288
hostnames	
IS-IS, displaying.....	363

I

- ignore-attached-bit statement.....240
- ignore-lsp-metrics statement
 - IS-IS.....241
- IGP
 - shortcuts
 - overview.....158
- inet statement.....290
- interface statement
 - IS-IS.....242
- interfaces
 - IS-IS, verifying.....535
- IP addresses
 - router mapping.....41
- ipv4-multicast statement.....244
 - usage guidelines.....110
- ipv4-multicast-metric statement.....245
 - usage guidelines.....110
- IPv6
 - IS-IS.....126
- ipv6-multicast statement.....245
- ipv6-multicast-metric statement.....246
- ipv6-unicast statement.....247
 - usage guidelines.....132
- ipv6-unicast-metric statement.....248
- IS-IS
 - addresses.....4
 - adjacency database entries, clearing.....307
 - authentication.....43, 45, 216, 266
 - CSNP.....267
 - hello.....268
 - hitless keychain.....46
 - PSNP.....272
 - authentication keychain.....217, 233
 - authentication, displaying.....339, 524
 - backup coverage
 - displaying.....341
 - backup MPLS LSPs.....343
 - backup paths
 - SPF calculations.....345
 - backup-spf-options statement.....219
 - BFD.....91, 220
 - CLNS.....222, 223
 - complete sequence number
 - PDUs.....6, 195, 225, 293
 - configuring on logical systems.....30, 82
 - designated router.....281
 - disabling.....226
 - IPv4 multicast topology.....110
 - IPv4 unicast topology.....110
 - IPv6 multicast topology.....110
 - IPv6 unicast topology.....132
 - dual stacking.....126
 - dynamic overload bit, resetting.....311
 - enabling.....249
 - errored packets.....293
 - graceful restart.....231
 - hello
 - interval.....235
 - packet authentication.....234
 - packet authentication key.....232
 - PDUs.....6, 293
 - hold time.....238
 - hold-down timer
 - disabling.....266
 - hostname database, displaying.....363
 - interfaces.....242
 - interfaces, displaying.....365
 - IP fast reroute.....141, 145
 - IPv4 unicast topology.....273
 - IPv6 unicast topology.....247, 272
 - label-switched path.....250
 - LDP synchronization.....185, 227
 - hold time.....239
 - level properties
 - global.....252
 - interfaces.....253
 - link protection.....143
 - link-protection statement.....254
 - link-state database entries
 - clearing.....309
 - displaying.....351
 - link-state PDUs.....6, 293
 - interval.....190, 256
 - lifetime.....257
 - tracing.....293
 - loop-free alternate routes.....141, 145
 - loose authentication.....45, 254
 - mesh groups.....200, 261
 - metrics.....284
 - IPv6.....248
 - multicast.....245, 246
 - normal.....262
 - traffic engineering.....291
 - wide.....180, 181, 299

multicast topologies.....	109, 110, 244, 245
IPv4.....	268
IPv6.....	270
multilevel.....	20
neighbors, displaying.....	335
network PDUs.....	3
no-eligible-backup statement.....	267
node link protection.....	144, 273
NSAP.....	4
overloaded, marking router as.....	274
padding.....	236
partial sequence number PDUs.....	6, 293
PDUs.....	6
point-to-point interface.....	278
policy, routing.....	73, 228
preferences.....	73, 229, 279
prefix limit.....	280
redistributing OSPF routes into.....	64
route tagging.....	88
routes, displaying.....	372
routing domains.....	253
RSVP LSP backup paths.....	144
SPF calculations, displaying.....	376
SPF delay calculations.....	293
supported software standards.....	9
topology.....	292
tracing operations.....	293
traffic engineering support.....	159, 226, 287, 296
traffic statistics	
clearing.....	313
displaying.....	381
wide metrics.....	180, 181
IS-IS (Intermediate System-to-Intermediate System)	
about designated routers.....	39
BFD authentication.....	104
checksum.....	51
configuring.....	14
configuring designated router election	
priority.....	40
for CLNS route exchange.....	207
IS-IS protocol.....	544
adjacencies	
status.....	547
verifying	535
broken network topology, figure	533
configuration	
Level 1 router	544
Level 1/Level 2 router	542
Level 2 router	546
configuration, verifying.....	536
detail statement.....	548, 551
displaying details.....	548
flag statement.....	548, 551, 552
hello statement.....	552
interfaces, verifying	535
link-state PDUs	
analyzing in detail.....	550
lsp statement.....	551
network topology	
detailed figure.....	542
figure	548
levels, figure	541
receive statement.....	552
run show isis interface command.....	543
run show log command.....	549, 553
send statement.....	552
set flag command.....	548, 551, 552
show configuration protocols isis	
command.....	536
show isis adjacency command.....	535
show isis interface command	535
show mpls lsp extensive command.....	534, 538
trace messages.....	548
tracing	
configuring.....	552
flags, table.....	549
verifying.....	533
isis statement.....	249
ISO	
addresses.....	4
system identifier.....	4
K	
keychain	
IS-IS.....	45
L	
label-switched-path statement	
IS-IS.....	250
usage guidelines.....	173
ldp-synchronization statement	
IS-IS.....	251
usage guidelines.....	185

level statement	
IS-IS	
interfaces.....	253
protocol.....	252
usage guidelines.....	20
link-protection statement	
IS-IS.....	143, 254
link-state PDUs.....	6
link-state PDUs, analyzing IS-IS.....	550
logical systems	
configuring default route.....	82
configuring IS-IS.....	30, 82
configuring routing policy.....	82
overview.....	28
loose-authentication-check statement	
IS-IS.....	254
usage guidelines.....	45
lsp	
IS-IS protocol	
statement	551
tracing flag.....	550
lsp (tracing flag).....	293
lsp-equal-cost statement.....	255
lsp-generation (tracing flag).....	293
lsp-generation, IS-IS protocol tracing flag.....	550
lsp-interval statement.....	256
usage guidelines.....	190
lsp-lifetime statement.....	257
LSPs	
IGP shortcuts.....	158
M	
manuals	
comments on.....	xix
max-areas statement.....	258
max-hello-size statement.....	259
max-lsp-size statement.....	259
max-snp-size statement.....	260
mesh groups.....	200, 261
mesh-group statement.....	261
usage guidelines.....	200
metric statement	
IS-IS.....	262
metrics	
IS-IS.....	284
minimum-interval statement	
IS-IS.....	220
minimum-receive-interval statement	
IS-IS.....	220
usage guidelines.....	91
multicast-rpf-routes statement.....	263
multilevel IS-IS.....	20
multipath statement	
IS-IS.....	264
multiplier statement	
IS-IS.....	220
usage guidelines.....	91
N	
names	
names.....	41
network PDUs.....	3
network service access point.....	4
network topology	
IS-IS protocol, detailed figure.....	542
IS-IS protocol, figure.....	547
levels IS-IS protocol, figure.....	541
next hops	
routes sent to, displaying.....	476
no-adaptation statement	
BFD (IS-IS)	
usage guidelines.....	91
IS-IS.....	220
no-adjacency-down-notification statement	
IS-IS.....	265
no-adjacency-holddown statement.....	266
no-authentication-check statement.....	266
usage guidelines.....	43
no-csnnp-authentication statement.....	267
usage guidelines.....	43
no-eligible-backup statement	
IS-IS.....	267
no-hello-authentication statement.....	268
usage guidelines.....	43
no-ipv4-multicast statement.....	268
no-ipv4-routing statement.....	269
no-ipv6-multicast statement.....	270
no-ipv6-routing statement.....	271
no-ipv6-unicast statement.....	272
no-psnp-authentication statement.....	272
usage guidelines.....	43
no-unicast-topology statement.....	273
node-link-protection statement.....	273
usage guidelines	
IS-IS.....	144
NPDUs.....	3

NSAP.....	4	receive statements	
O		IS-IS protocol.....	552
OSPF		redistributing routes	
redistributing routes into IS-IS.....	64	OSPF into IS-IS.....	64
overload bit, resetting for IS-IS.....	311	reference-bandwidth statement	
overload statement		IS-IS.....	284
IS-IS.....	274	restart command.....	317
P		restart-duration statement	
packets		IS-IS.....	231
IS-IS protocol tracing flag.....	550	restarting	
received.....	552	software processes.....	317
sent.....	552	RFC 3719, Recommendations for Interoperable	
packets (tracing flag)		Networks using Intermediate System to	
IS-IS.....	293	Intermediate System (IS-IS).....	9
parentheses, in syntax descriptions.....	xix	rib-group statement	
partial sequence number PDUs.....	6	IS-IS.....	285
passive statement		route advertisements, displaying.....	398
IS-IS.....	277	route, displaying	
PDUs.....	6	next-hop.....	476
ping clns command.....	315	routers	
point-to-point statement.....	278	configuration	
policy, routing		IS-IS protocol	544
communities.....	72	Level 1 IS-IS, configuring.....	544
framework.....	61	Level 1/Level 2 IS-IS, configuring.....	542
IS-IS.....	73, 228	Level 2 IS-IS, configuring.....	546
overview.....	61	names	
preference statement		configuring.....	41
IS-IS.....	279	mapping to IP addresses.....	41
preferences		routes, displaying	
IS-IS.....	229, 279	active.....	388
prefix limit		active path.....	393
IS-IS.....	280	advertising protocol.....	398
prefix-export-limit statement		all.....	403
IS-IS.....	280	best.....	405
priority statement		brief information.....	408
IS-IS.....	281	detailed information.....	410
processes		extensive information.....	432
restarting.....	317	hidden.....	463
protocols statement.....	282	in a specific routing table.....	507
psn (tracing flag).....	293	in the forwarding table.....	449
psn, IS-IS protocol tracing flag.....	550	inactive path.....	466
PSNP IS-IS.....	6	instances.....	469
R		learned from a specific protocol.....	487
rapid-runs statement		matching the specified address.....	427
IS-IS.....	288	policy-based route export.....	429
		received through a neighbor.....	499
		sent to a specific interface.....	482
		terse information.....	521

routing policies
 BGP export, for CLNS.....207
 displaying.....384
 testing the configuration for.....527
 routing policy See policy, routing
 configuring on logical systems.....82
 routing protocol process memory
 FAQ.....555, 556
 routing-instances statement.....286
 run show isis interface command.....543, 544, 546
 run show log command
 IS-IS protocol549, 551, 553

S

send statement
 IS-IS protocol.....552
 set flag command
 IS-IS protocol548, 551, 552
 set flag hello detail command548
 set flag hello send command552
 set flag lsp detail command551
 shortcuts statement
 IS-IS.....287
 usage guidelines.....159
 shortest path first.....376
 show bfd session command.....327
 show configuration protocols isis command536
 show isis adjacency brief command.....18
 show isis adjacency command.....335, 535
 show isis adjacency extensive command.....18
 explanation.....19
 show isis authentication command.....339
 show isis backup coverage.....341
 show isis backup spf results.....345
 show isis context-identifier command.....349
 show isis database command.....351
 show isis hostname command.....363
 show isis interface command.....204, 365, 535
 show isis overview command.....369
 show isis route command.....372
 show isis spf command.....376
 show isis statistics command.....205, 381
 show isis-backup label-switched-path
 command.....343
 show log command.....53
 show mpls lsp extensive command
 IS-IS protocol534, 538
 show policy command.....384
 show policy conditions command.....386
 show route active-path command.....393
 show route advertising-protocol command.....398
 show route all command.....403
 show route best command.....405
 show route brief command.....408
 show route command.....388
 show route detail command.....410
 show route exact command.....427
 show route export command.....429
 show route extensive command.....432
 show route forwarding-table command.....449
 show route hidden command.....463
 show route inactive-path command.....466
 show route instance command.....469
 show route next-hop command.....476
 show route output command.....482
 show route protocol command.....487
 show route receive-protocol command.....499
 show route table command.....507
 show route terse command.....521
 show security keychain command.....524
 show bfd session extensive command.....106
 show isis interface brief command.....17
 show isis interface detail command.....17
 explanation.....17
 spf
 IS-IS protocol tracing flag.....550
 spf (tracing flag)
 IS-IS.....293
 spf-options statement
 IS-IS.....288
 static-host-mapping statement.....290
 usage guidelines.....40
 support, technical See technical support
 syntax conventions.....xviii
 sysid statement.....290
 usage guidelines.....40
 system identifier.....4
 system identifier, IS-IS
 configuring.....40

T

te-metric statement
 IS-IS.....291
 technical support
 contacting JTAC.....xx
 test policy command.....527

threshold statement	
IS-IS.....	220
usage guidelines.....	91
topologies statement.....	292
traceoptions statement	
IS-IS.....	293
traceroute clns command.....	529
tracing flags	
csn.....	293
error	
IS-IS.....	293
graceful restart	
IS-IS.....	293
hello	
IS-IS.....	293
IS-IS protocol, table	549
lsp.....	293
lsp-generation.....	293
packets	
IS-IS.....	293
psn.....	293
spf	
IS-IS.....	293
tracing operations	
IS-IS.....	293
traffic engineering	
IGP shortcuts.....	158
traffic-engineering statement	
IS-IS.....	296
usage guidelines.....	159
transmit-interval statement	
IS-IS.....	220
 V	
verification	
BFD for IS-IS.....	97
IS-IS policy.....	86
multicast topology for IS-IS.....	115
version statement	
IS-IS.....	220
usage guidelines.....	91
 W	
wide-metrics-only statement.....	299
usage guidelines.....	180, 181