



---

# IS-IS Feature Guide for NFX Series Devices



---

Modified: 2018-07-08



Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be 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.

*IS-IS Feature Guide for NFX Series Devices*

Copyright © 2018 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 <https://www.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.



# Table of Contents

|                  |  |           |
|------------------|--|-----------|
|                  | About the Documentation . . . . .  | xi        |
|                  | Documentation and Release Notes . . . . .  | xi        |
|                  | Using the Examples in This Manual . . . . .  | xi        |
|                  | Merging a Full Example . . . . .   | xii       |
|                  | Merging a Snippet . . . . .  | xii       |
|                  | Documentation Conventions . . . . .  | xiii      |
|                  | Documentation Feedback . . . . .   | xv        |
|                  | Requesting Technical Support . . . . .   | xv        |
|                  | Self-Help Online Tools and Resources . . . . .   | xv        |
|                  | Opening a Case with JTAC . . . . .   | xvi       |
| <b>Part 1</b>    | <b>Overview</b>  |           |
| <b>Chapter 1</b> | <b>IS-IS . . . . .</b>   | <b>3</b>  |
|                  | IS-IS Overview . . . . .   | 3         |
|                  | IS-IS Terminology . . . . .  | 4         |
|                  | 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<br>Both IS-IS Levels . . . . . | 7         |
| <b>Part 2</b>    | <b>Configuring IS-IS Features</b>  |           |
| <b>Chapter 2</b> | <b>Configuring IS-IS . . . . .</b>   | <b>11</b> |
|                  | Example: Configuring IS-IS . . . . .   | 11        |
|                  | Example: Configuring a Multi-Level IS-IS Topology to Control Interarea<br>Flooding . . . . .             | 17        |
|                  | Configuring IS-IS Authentication . . . . .   | 26        |
|                  | Configuring IS-IS Authentication Without Network-Wide Deployment . . . . .                               | 27        |
|                  | Example: Redistributing OSPF Routes into IS-IS . . . . .   | 28        |
|                  | Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies . . . . .                                    | 36        |
|                  | Understanding IS-IS IPv4 and IPv6 Unicast Topologies . . . . .   | 37        |
|                  | Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies . . . . .                                    | 37        |
|                  | Example: Configuring IS-IS Multicast Topology . . . . .  | 46        |
|                  | IS-IS Multicast Topologies Overview . . . . .  | 46        |
|                  | Example: Configuring IS-IS Multicast Topology . . . . .  | 47        |
|                  | Example: Configuring an IS-IS Default Route Policy on Logical Systems . . . . .                          | 62        |



|                  |   |            |
|------------------|---|------------|
|                  | Example: Configuring IS-IS for CLNS . . . . .   | 68         |
|                  | Understanding IS-IS for CLNS . . . . .  | 68         |
|                  | Example: Configuring IS-IS for CLNS . . . . .   | 69         |
|                  | Example: Configuring IS-IS Designated Routers . . . . .                               | 71         |
|                  | Understanding IS-IS Designated Routers . . . . .                                      | 71         |
|                  | Configuring Designated Router Election Priority for IS-IS . . . . .                   | 71         |
|                  | Example: Enabling Packet Checksums on IS-IS Interfaces for Error Checking . . . . .   | 72         |
| <b>Chapter 3</b> | <b>Configuring BFD Authentication for IS-IS . . . . .</b>                             | <b>75</b>  |
|                  | Understanding BFD Authentication for IS-IS . . . . .                                  | 75         |
|                  | BFD Authentication Algorithms . . . . .   | 75         |
|                  | Security Authentication Keychains . . . . .   | 76         |
|                  | Strict Versus Loose Authentication . . . . .  | 76         |
|                  | Example: Configuring BFD for IS-IS . . . . .  | 77         |
|                  | Example: Configuring BFD Authentication for IS-IS . . . . .                           | 83         |
| <b>Chapter 4</b> | <b>Configuring Hitless Authentication Key Rollover for IS-IS . . . . .</b>            | <b>87</b>  |
|                  | Understanding Hitless Authentication Key Rollover for IS-IS . . . . .                 | 87         |
|                  | Example: Configuring Hitless Authentication Key Rollover for IS-IS . . . . .          | 88         |
| <b>Chapter 5</b> | <b>Configuring Loop-Free Alternate Routes for IS-IS . . . . .</b>                     | <b>95</b>  |
|                  | Example: Configuring Link and Node Protection for IS-IS Routes . . . . .              | 95         |
|                  | Understanding Loop-Free Alternate Routes for IS-IS . . . . .                          | 95         |
|                  | Configuring Link Protection for IS-IS . . . . .                                       | 97         |
|                  | Configuring Node-Link Protection for IS-IS . . . . .                                  | 97         |
|                  | Excluding an IS-IS Interface as a Backup for Protected Interfaces . . . . .           | 98         |
|                  | Configuring RSVP Label-Switched Paths as Backup Paths for IS-IS . . . . .             | 98         |
|                  | Using Operational Mode Commands to Monitor Protected IS-IS Routes . . . . .           | 99         |
|                  | Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN . . . . . | 99         |
| <b>Chapter 6</b> | <b>Configuring Segment Routing . . . . .</b>  | <b>111</b> |
|                  | Understanding Source Packet Routing in Networking (SPRING) . . . . .                  | 111        |
|                  | Support for SRGB, Anycast, and Prefix Segments in SPRING for IS-IS Protocol . . . . . | 114        |
|                  | Configuring Anycast and Prefix segments in SPRING for IS-IS Protocol . . . . .        | 116        |
| <b>Part 3</b>    | <b>Configuration Statements and Operational Commands</b>                              |            |
| <b>Chapter 7</b> | <b>Configuration Statements . . . . .</b>   | <b>123</b> |
|                  | adjacency-segment . . . . .   | 125        |
|                  | authentication-key (Protocols IS-IS) . . . . .  | 126        |
|                  | authentication-key-chain (Protocols IS-IS) . . . . .                                  | 127        |
|                  | authentication-type (Protocols IS-IS) . . . . .                                       | 128        |
|                  | bfd-liveness-detection (Protocols IS-IS) . . . . .                                    | 129        |
|                  | checksum (Protocols IS-IS) . . . . .  | 132        |
|                  | csnp-interval . . . . .   | 133        |
|                  | disable (Protocols IS-IS) . . . . .   | 134        |
|                  | export . . . . .  | 135        |
|                  | external-preference (Protocols IS-IS) . . . . .                                       | 136        |



|  |     |
|--|-----|
| family (Protocols IS-IS) .....                         | 137 |
| graceful-restart (Protocols IS-IS) .....               | 138 |
| hello-authentication-key .....                         | 139 |
| hello-authentication-key-chain .....                   | 140 |
| hello-authentication-type .....                        | 141 |
| hello-interval (Protocols IS-IS) .....                 | 142 |
| hello-padding .....                                    | 144 |
| hold-time (Protocols IS-IS) .....                      | 146 |
| ignore-attached-bit .....                              | 148 |
| interface (Protocols IS-IS) .....                      | 149 |
| ipv4-multicast .....                                   | 151 |
| ipv4-multicast-metric .....                            | 152 |
| ipv6-multicast .....                                   | 153 |
| ipv6-multicast-metric .....                            | 154 |
| ipv6-unicast .....                                     | 155 |
| ipv6-unicast-metric .....                              | 156 |
| isis .....   | 157 |
| level (Global IS-IS) .....                             | 158 |
| link-protection (Protocols IS-IS) .....                | 159 |
| loose-authentication-check .....                       | 160 |
| lsp-interval .....                                     | 161 |
| lsp-lifetime .....                                     | 163 |
| max-areas .....  | 164 |
| mesh-group (Protocols IS-IS) .....                     | 165 |
| metric (Protocols IS-IS) .....                         | 166 |
| no-adjacency-holddown .....                            | 167 |
| no-advertise-adjacency-segment (Protocols IS-IS) ..... | 168 |
| no-authentication-check .....                          | 169 |
| no-csnp-authentication .....                           | 170 |
| no-eligible-backup (Protocols IS-IS) .....             | 171 |
| no-hello-authentication .....                          | 172 |
| no-ipv4-multicast .....                                | 173 |
| no-ipv4-routing .....                                  | 174 |
| no-ipv6-multicast .....                                | 175 |
| no-ipv6-routing .....                                  | 176 |
| no-ipv6-unicast .....                                  | 177 |
| no-psnp-authentication .....                           | 178 |
| no-unicast-topology .....                              | 179 |
| node-link-protection (Protocols IS-IS) .....           | 180 |
| node-segment (Protocols IS-IS) .....                   | 181 |
| overload (Protocols IS-IS) .....                       | 183 |
| passive (Protocols IS-IS) .....                        | 186 |
| point-to-point .....                                   | 188 |
| preference (Protocols IS-IS) .....                     | 189 |
| prefix-export-limit (Protocols IS-IS) .....            | 190 |
| prefix-segment .....                                   | 191 |
| priority (Protocols IS-IS) .....                       | 192 |
| reference-bandwidth (Protocols IS-IS) .....            | 193 |
| rib-group (Protocols IS-IS) .....                      | 194 |



|                  |  |            |
|------------------|--|------------|
|                  | source-packet-routing (Protocols IS-IS and OSPF) ..... | 195        |
|                  | source-packet-routing-disable (Protocols IS-IS) .....  | 197        |
|                  | spf-options (Protocols IS-IS) .....                    | 198        |
|                  | srgb .....   | 200        |
|                  | topologies (Protocols IS-IS) .....                     | 201        |
|                  | traceoptions (Protocols IS-IS) .....                   | 202        |
|                  | traffic-engineering (Protocols IS-IS) .....            | 205        |
|                  | use-source-packet-routing (Protocols IS-IS) .....      | 209        |
|                  | wide-metrics-only .....                                | 210        |
| <b>Chapter 8</b> | <b>Operational Commands .....</b>                      | <b>211</b> |
|                  | clear isis adjacency .....                             | 212        |
|                  | clear isis database .....                              | 214        |
|                  | clear isis overload .....                              | 216        |
|                  | clear isis statistics .....                            | 218        |
|                  | show isis adjacency .....                              | 220        |
|                  | show isis authentication .....                         | 225        |
|                  | show isis backup coverage .....                        | 227        |
|                  | show isis backup label-switched-path .....             | 229        |
|                  | show isis backup spf results .....                     | 231        |
|                  | show isis database .....                               | 235        |
|                  | show isis hostname .....                               | 250        |
|                  | show isis interface .....                              | 252        |
|                  | show isis overview .....                               | 257        |
|                  | show isis route .....                                  | 261        |
|                  | show isis spf .....                                    | 265        |
|                  | show isis statistics .....                             | 270        |



# List of Figures

|                  |   |           |
|------------------|---|-----------|
| <b>Part 1</b>    | <b>Overview</b>   |           |
| <b>Chapter 1</b> | <b>IS-IS</b>  | <b>3</b>  |
|                  | Figure 1: Install Default Route to Nearest Routing Device That Operates at Both Level 1 and Level 2 | 7         |
| <b>Part 2</b>    | <b>Configuring IS-IS Features</b>   |           |
| <b>Chapter 2</b> | <b>Configuring IS-IS</b>  | <b>11</b> |
|                  | Figure 2: Simple IS-IS Topology   | 12        |
|                  | Figure 3: IS-IS Multi-Level Topology  | 18        |
|                  | Figure 4: IS-IS Route Redistribution Topology   | 29        |
|                  | Figure 5: IS-IS IPv4 and IPv6 Unicast Topologies  | 39        |
|                  | Figure 6: Configuring IS-IS Multicast Topology  | 48        |
|                  | Figure 7: IS-IS with a Default Route to an ISP  | 63        |
|                  | Figure 8: IS-IS Checksum Topology   | 73        |
| <b>Chapter 3</b> | <b>Configuring BFD Authentication for IS-IS</b>   | <b>75</b> |
|                  | Figure 9: Configuring BFD for IS-IS   | 77        |
|                  | Figure 10: IS-IS BFD Authentication Topology  | 83        |
| <b>Chapter 4</b> | <b>Configuring Hitless Authentication Key Rollover for IS-IS</b>                                    | <b>87</b> |
|                  | Figure 11: Hitless Authentication Key Rollover for IS-IS  | 89        |
| <b>Chapter 5</b> | <b>Configuring Loop-Free Alternate Routes for IS-IS</b>   | <b>95</b> |
|                  | Figure 12: Link Protection and Node-Link Protection Comparison for IS-IS Routes                     | 96        |
|                  | Figure 13: IS-IS Node-Link Protection Topology  | 100       |







# List of Tables

|                  |   |            |
|------------------|---|------------|
|                  | <b>About the Documentation</b> . . . . .                                | <b>xi</b>  |
|                  | Table 1: Notice Icons . . . . .   | xiii       |
|                  | Table 2: Text and Syntax Conventions . . . . .                          | xiii       |
| <b>Part 2</b>    | <b>Configuring IS-IS Features</b>                                       |            |
| <b>Chapter 2</b> | <b>Configuring IS-IS</b> . . . . .                                      | <b>11</b>  |
|                  | Table 3: IPv4 Statements . . . . .                                      | 46         |
|                  | Table 4: IPv6 Statements . . . . .                                      | 47         |
| <b>Part 3</b>    | <b>Configuration Statements and Operational Commands</b>                |            |
| <b>Chapter 7</b> | <b>Configuration Statements</b> . . . . .                               | <b>123</b> |
|                  | Table 5: Default Metric Values for Routes Exported into IS-IS . . . . . | 166        |
| <b>Chapter 8</b> | <b>Operational Commands</b> . . . . .                                   | <b>211</b> |
|                  | Table 6: show isis adjacency Output Fields . . . . .                    | 221        |
|                  | Table 7: show isis authentication Output Fields . . . . .               | 225        |
|                  | Table 8: show isis backup coverage Output Fields . . . . .              | 227        |
|                  | Table 9: show isis backup label-switched-path Output Fields . . . . .   | 229        |
|                  | Table 10: show isis backup spf results Output Fields . . . . .          | 232        |
|                  | Table 11: show isis database Output Fields . . . . .                    | 236        |
|                  | Table 12: show isis hostname Output Fields . . . . .                    | 250        |
|                  | Table 13: show isis interface Output Fields . . . . .                   | 253        |
|                  | Table 14: show isis overview Output Fields . . . . .                    | 257        |
|                  | Table 15: show isis route Output Fields . . . . .                       | 262        |
|                  | Table 16: show isis spf Output Fields . . . . .                         | 266        |
|                  | Table 17: show isis statistics Output Fields . . . . .                  | 271        |







# About the Documentation

- Documentation and Release Notes on page xi
- Using the Examples in This Manual on page xi
- Documentation Conventions on page xiii
- Documentation Feedback on page xv
- Requesting Technical Support on page xv

## 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 <https://www.juniper.net/documentation/>.

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 <https://www.juniper.net/books>.

## 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 [CLI Explorer](#).

## Documentation Conventions

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

Table 2: Text and Syntax Conventions

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



Table 2: Text and Syntax Conventions (continued)

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



Table 2: Text and Syntax Conventions (continued)

| Convention                   | Description   | Examples   |
|------------------------------|---|--|
| > (bold right angle bracket) | Separates levels in a hierarchy of menu selections. | In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> . |

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <https://www.juniper.net/documentation/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <https://www.juniper.net/documentation/feedback/>.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software 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 Partner Support Service 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 <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://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: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>



- Download the latest versions of software and review release notes:  
<https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications:  
<https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum:  
<https://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

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

## 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 <https://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 <https://www.juniper.net/support/requesting-support.html>.



## PART 1

# Overview

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







## CHAPTER 1

# 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.



**NOTE:** See *Platforms/FPCs That Cannot Forward TCC Encapsulated ISO Traffic* to find a list of those devices and FPC configurations that cannot pass ISO traffic when encapsulated in TCC format.

This section discusses the following topics:

- [IS-IS Terminology on page 4](#)
- [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<sup>®</sup> 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**
- *Understanding IS-IS Configuration*
  - [Example: Configuring IS-IS on page 11](#)



## PART 2

# Configuring IS-IS Features

- [Configuring IS-IS on page 11](#)
- [Configuring BFD Authentication for IS-IS on page 75](#)
- [Configuring Hitless Authentication Key Rollover for IS-IS on page 87](#)
- [Configuring Loop-Free Alternate Routes for IS-IS on page 95](#)
- [Configuring Segment Routing on page 111](#)







## CHAPTER 2

# Configuring IS-IS

- [Example: Configuring IS-IS on page 11](#)
- [Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 17](#)
- [Configuring IS-IS Authentication on page 26](#)
- [Configuring IS-IS Authentication Without Network-Wide Deployment on page 27](#)
- [Example: Redistributing OSPF Routes into IS-IS on page 28](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 36](#)
- [Example: Configuring IS-IS Multicast Topology on page 46](#)
- [Example: Configuring an IS-IS Default Route Policy on Logical Systems on page 62](#)
- [Example: Configuring IS-IS for CLNS on page 68](#)
- [Example: Configuring IS-IS Designated Routers on page 71](#)
- [Example: Enabling Packet Checksums on IS-IS Interfaces for Error Checking on page 72](#)

## Example: Configuring IS-IS

---

This example shows how to configure IS-IS.

- [Requirements on page 11](#)
- [Overview on page 11](#)
- [Configuration on page 12](#)
- [Verification on page 14](#)

### 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 12 shows the topology used in this example.

Figure 2: Simple IS-IS Topology



“CLI Quick Configuration” on page 12 shows the configuration for both of the devices in Figure 2 on page 12. The section “Step-by-Step Procedure” on page 12 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 14](#)
- [Verifying IS-IS Interface Configuration in Detail on page 14](#)
- [Verifying IS-IS Adjacencies on page 15](#)
- [Verifying IS-IS Adjacencies in Detail on page 16](#)

### 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              3  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        Seenself
```

**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**
- [Understanding IS-IS Configuration](#)
  - [Example: Configuring IS-IS for GRES with Graceful Restart](#)
  - [Configuring Designated Router Election Priority for IS-IS on page 71](#)
  - [Junos OS Feature Support Reference for SRX Series and J Series Devices](#)

---

## Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding

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

- [Requirements on page 17](#)
- [Overview on page 17](#)
- [Configuration on page 18](#)
- [Verification on page 22](#)

### 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 18](#) shows the topology used in this example.



Figure 3: IS-IS Multi-Level Topology



“CLI Quick Configuration” on page 18 shows the configuration for all of the devices in Figure 3 on page 18. The section “Step-by-Step Procedure” on page 20 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 23](#)
- [Verifying IS-IS Adjacencies on page 23](#)
- [Examining the IS-IS Database on page 24](#)



### 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
100.0              3  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 Level 1/Level 2 (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**

- *Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups*



## 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 IS-IS Authentication Without Network-Wide Deployment on page 27](#)

## Configuring IS-IS 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

- [Understanding Hitless Authentication Key Rollover for IS-IS on page 87](#)
- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 88](#)



## Example: Redistributing OSPF Routes into IS-IS

---

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

- [Requirements on page 28](#)
- [Overview on page 28](#)
- [Configuration on page 29](#)
- [Verification on page 34](#)

### 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 4 on page 29](#) shows the topology used in this example.



Figure 4: IS-IS Route Redistribution Topology



“CLI Quick Configuration” on page 29 shows the configuration for all of the devices in Figure 4 on page 29. The section “Step-by-Step Procedure” on page 30 describes the steps on Device R2. “Step-by-Step Procedure” on page 31 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-1/2/0.0
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 35](#)
- [Verifying Route Redistribution on page 35](#)
- [Verifying Connectivity on page 36](#)



### 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**

- [Understanding Routing Policies](#)

---

## Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

---

- [Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 37](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 37](#)



## 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.

**See Also** • [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 37](#)

## 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 37](#)
- [Overview on page 37](#)
- [Configuration on page 39](#)
- [Verification on page 43](#)

---

### 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 5 on page 39](#) 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 5: IS-IS IPv4 and IPv6 Unicast Topologies



“CLI Quick Configuration” on page 39 shows the configuration for all of the devices in Figure 5 on page 39. The section “Step-by-Step Procedure” on page 41 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 43](#)
- [Checking the IS-IS SPF Calculations on page 44](#)
- [Checking the Tcpdump Output on page 45](#)

#### *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.

**See Also** • *Example: Configuring IS-IS Dual Stacking of IPv4 and IPv6 Unicast Addresses*



- Related Documentation**
- [Example: Configuring IS-IS Dual Stacking of IPv4 and IPv6 Unicast Addresses](#)
  - [Understanding Dual Stacking of IPv4 and IPv6 Unicast Addresses](#)

## Example: Configuring IS-IS Multicast Topology

- [IS-IS Multicast Topologies Overview on page 46](#)
- [Example: Configuring IS-IS Multicast Topology on page 47](#)

### 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`.



**NOTE:** For the multicast metric commands, these are interface-specific settings, not global.

Table 3 on page 46 lists the various IPv4 statements you can use to configure IS-IS topologies.

**Table 3: IPv4 Statements**

| Statement                   | Description                                   |
|-----------------------------|---|
| <code>ipv4-multicast</code> | Enables an alternate IPv4 multicast topology. |



*Table 3: IPv4 Statements (continued)*

| Statement  | Description   |
|--|---|
| <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 4 on page 47](#) lists the various IPv6 statements you can use to configure IS-IS topologies.

*Table 4: IPv6 Statements*

| Statement  | Description   |
|--|---|
| <code>ipv6-multicast</code>                      | Enables an alternate IPv6 multicast topology.                             |
| <code>ipv6-unicast</code>                        | Enables an alternate IPv6 unicast 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.

**See Also** • [Example: Configuring IS-IS Multicast Topology on page 47](#)

## Example: Configuring IS-IS Multicast Topology

This example shows how to configure a multicast topology for an IS-IS network.

- [Requirements on page 47](#)
- [Overview on page 48](#)
- [Configuration on page 48](#)
- [Verification on page 52](#)

### Requirements

Before you begin, configure IS-IS on all routers. See [“Example: Configuring IS-IS” on page 11](#) 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 6 on page 48 shows the sample network.

Figure 6: 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 52](#)
- [Verifying That IS-IS Is Configured on page 54](#)
- [Verifying the Configured Multicast Metric Values on page 56](#)
- [Verifying the Configuration of the Multicast Topology on page 57](#)

#### *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.

**See Also**

- [IS-IS Multicast Topologies Overview on page 46](#)

**Related Documentation**

- [Understanding Multitopology Routing in Conjunction with PIM](#)
- [Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 36](#)

---

## 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 62](#)
- [Overview on page 62](#)
- [Configuration on page 63](#)
- [Verification on page 66](#)

### 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 7 on page 63](#) shows the sample network.



Figure 7: IS-IS with a Default Route to an ISP



g040918

## 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, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter commit from configuration mode.

```

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 **IS-IS**. 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: Creating an Interface on a Logical System](#)

---

## Example: Configuring IS-IS for CLNS

- [Understanding IS-IS for CLNS on page 68](#)
- [Example: Configuring IS-IS for CLNS on page 69](#)

### 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.

**See Also**

- [CLNS Overview](#)
- [Example: Configuring IS-IS for CLNS on page 69](#)



## Example: Configuring IS-IS for CLNS

This example shows how to create a routing instance and enable IS-IS protocol on all interfaces.

- [Requirements on page 69](#)
- [Overview on page 69](#)
- [Configuration on page 69](#)
- [Verification on page 70](#)

### Requirements

Before you begin, configure the network interfaces. See [Interfaces Feature 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, and define BGP export policy name (`dist-bgp`), family (`ISO`), and protocol (`BP`), 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, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
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 instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure IS-IS for CLNS:

1. Configure the routing instance.  

```
[edit]
user@host# edit routing-instances aaaa
```
2. Enable CLNS routing.  

```
[edit routing-instances aaaa]
user@host# set protocols isis clns-routing
```



3. Enable IS-IS on all interfaces.  

```
[edit routing-instances aaaa]  
user@host# set protocols isis interface all
```
4. (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
```
5. Define the BGP export policy name, family, and protocol.  

```
[edit policy-options]  
user@host# set policy-statement dist-bgp from family iso protocol bgp
```
6. Define the action for the export policy.  

```
[edit policy-options]  
user@host# set policy-statement dist-bgp then accept
```
7. 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** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]  
user@host# show routing-instances  
aaaa {  
  protocols {  
    isis {  
      export dist-bgp;  
      no-ipv4-routing;  
      no-ipv6-routing;  
      clns-routing;  
      interface all;  
    }  
  }  
}
```

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

---

### Verification

Confirm that the configuration is working properly.

- [Verifying Routing-Instance for CLNS on page 70](#)
- [Verifying IS-IS for CLNS on page 71](#)

#### *Verifying Routing-Instance for CLNS*

**Purpose** Verify that the policy options are enabled for the routing instance.



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

#### *Verifying IS-IS for CLNS*

**Purpose** Verify that IS-IS is enabled.

**Action** From operational mode, enter the **show protocols** command.

- See Also**
- [CLNS Configuration Overview](#)
  - [Understanding IS-IS for CLNS on page 68](#)
  - [Verifying a CLNS VPN Configuration](#)

---

## Example: Configuring IS-IS Designated Routers

- [Understanding IS-IS Designated Routers on page 71](#)
- [Configuring Designated Router Election Priority for IS-IS on page 71](#)

### 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.

- See Also**
- [Junos OS Feature Support Reference for SRX Series and J Series Devices](#)
  - [Configuring Designated Router Election Priority for IS-IS on page 71](#)

### 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 11](#).

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
```

- See Also**
- [Understanding IS-IS Designated Routers on page 71](#)
  - [Example: Configuring IS-IS on page 11](#)

- Related Documentation**
- [Example: Configuring IS-IS on page 11](#)
  - [Understanding IS-IS Configuration](#)

---

## Example: Enabling Packet Checksums on IS-IS Interfaces for Error Checking

---

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

- [Requirements on page 72](#)
- [Overview on page 72](#)
- [Configuration on page 73](#)
- [Verification on page 74](#)

### Requirements

Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 11](#) 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 8 on page 73](#) shows the topology used in this example.



Figure 8: 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**

- *Understanding Checksums on IS-IS Interfaces for Error Checking*



## CHAPTER 3

# Configuring BFD Authentication for IS-IS

- [Understanding BFD Authentication for IS-IS on page 75](#)
- [Example: Configuring BFD for IS-IS on page 77](#)
- [Example: Configuring BFD Authentication for IS-IS on page 83](#)

## 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 75](#)
- [Security Authentication Keychains on page 76](#)
- [Strict Versus Loose Authentication on page 76](#)

## 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 I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method 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.

#### Related Documentation

- [Example: Configuring BFD Authentication for IS-IS on page 83](#)

## 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.



**NOTE:** BFD is not supported with ISIS for IPV6 on QFX10000 series switches.

- [Requirements on page 77](#)
- [Overview on page 77](#)
- [Configuration on page 77](#)
- [Verification on page 81](#)

### Requirements

Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 11](#) 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 9 on page 77](#) shows the sample network.

*Figure 9: 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 81](#)
- [Verifying That IS-IS Is Configured on page 81](#)
- [Verifying That BFD Is configured on page 82](#)

---

### 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 Time | Transmit Interval | Multiplier |
|----------|-------|-----------|-------------|-------------------|------------|
| 10.0.0.2 | Up    | so-0/0/0  | 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**

- [Understanding BFD for IS-IS](#)

## Example: Configuring BFD Authentication for IS-IS

This example shows how to configure BFD authentication for IS-IS.

- [Requirements on page 83](#)
- [Overview on page 83](#)
- [Configuration on page 84](#)
- [Verification on page 86](#)

### Requirements


Before you begin, configure IS-IS on both routers. See “[Example: Configuring IS-IS](#)” on [page 11](#) for information about the required IS-IS configuration.

### Overview

In this example, a BFD authentication keychain is configured with meticulous keyed MD5 authentication.

[Figure 10 on page 83](#) shows the topology used in this example.

**Figure 10: IS-IS BFD Authentication Topology**



“[CLI Quick Configuration](#)” on [page 84](#) shows the configuration for both of the devices in [Figure 10 on page 83](#). The section “[Step-by-Step Procedure](#)” on [page 84](#) 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 $ABC123
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 $ABC123
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 $ABC123
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 $ABC123
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 $ABC123
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 $ABC123
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 "$ABC123"
user@R1# set key 1 start-time "2012-5-31.13:00:00 -0700"
user@R1# set key 2 secret "$ABC123"
user@R1# set key 2 start-time "2013-5-31.13:00:00 -0700"

```



```
user@R1# set key 3 secret "$ABC123"
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 "$ABC123";
      start-time "2012-5-31.13:00:00 -0700";
    }
    key 2 {
      secret "$ABC123";
      start-time "2013-5-31.13:00:00 -0700";
    }
    key 3 {
      secret "$ABC123";
      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**

- *Configuring BFD Authentication for IS-IS*
- [Example: Configuring BFD for IS-IS on page 77](#)
- *Understanding BFD for IS-IS*



## CHAPTER 4

# Configuring Hitless Authentication Key Rollover for IS-IS

- [Understanding Hitless Authentication Key Rollover for IS-IS on page 87](#)
- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 88](#)

## 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.

### Related Documentation

- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 88](#)



## 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 88](#)
- [Overview on page 88](#)
- [Configuration on page 89](#)
- [Verification on page 92](#)

### 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 11 on page 89 shows the topology used in the example.

Figure 11: Hitless Authentication Key Rollover for IS-IS



This example shows the configuration for Router R0.

## Configuration

### CLI Quick Configuration

To quickly configure the hitless authentication key rollover for IS-IS, copy the following commands and paste the commands into the CLI.

```
[edit]
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:1F0C: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 "$ABC123"
set security authentication-key-chains key-chain base-key-global key 63 start-time
  "2011-8-6.06:54:00-0700"
set security authentication-key-chains key-chain base-key-global key 63 algorithm
  hmac-sha-1
set security authentication-key-chains key-chain base-key-global key 63 options
  isis-enhanced
set security authentication-key-chains key-chain base-key-inter key 0 secret "$ABC123"
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]
user@host# edit interfaces ge-0/0/0 unit 0
[edit interfaces ge-0/0/0 unit 0]
user@host# set description "interface A"
user@host# set family inet address 10.0.0.1/30
user@host# set family iso
user@host# set family inet6 address fe80::200:f8ff:fe21:67cf/128
user@host# exit
[edit]
user@host# edit interfaces ge-0/0/1 unit 0
[edit interfaces ge-0/0/1 unit 0]
user@host# set interfaces ge-0/0/1 unit 0 description "interface B"
user@host# set interfaces ge-0/0/1 unit 0 family inet address 10.0.0.5/30
user@host# set interfaces ge-0/0/1 unit 0 family iso
user@host# set interfaces ge-0/0/1 unit 0 family inet6 address
  10FB::C:ABC:1F0C:44DA/128
user@host# exit
[edit]
user@host# edit interfaces ge-0/0/2 unit 0
[edit interfaces ge-0/0/2 unit 0]
user@host# set description "interface C"
user@host# set family inet address 10.0.0.9/30
user@host# set interfaces ge-0/0/2 unit 0 family iso
user@host# set interfaces ge-0/0/2 unit 0 family inet6 address ff06::c3/128
user@host# exit

```

2. Configure one or more authentication keys.

```

[edit]
user@host# edit security authentication-key-chains key-chain base-key-global
[edit security authentication-key-chains key-chain base-key-global]
user@host# set key 63 secret "$ABC123"
user@host# set key 63 start-time "2011-8-6.06:54:00-0700"
user@host# set key 63 algorithm hmac-sha-1
user@host# set key 63 options isis-enhanced

```



```
user@host# exit
[edit]
user@host# edit security authentication-key-chains key-chain base-key-inter
[edit security authentication-key-chains key-chain base-key-inter]
user@host# set key 0 secret "$ABC123"
user@host# set key 0 start-time "2011-8-6.06:54:00-0700"
user@host# set key 0 algorithm md5
user@host# set key 0 options basic
user@host# exit
```

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

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

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

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

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

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

---

## Results

Confirm your configuration by entering the **show interfaces**, **show protocols**, and **show security** commands.

```
user@host# 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@host# 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@host# show security
authentication-key-chains {
    key-chain base-key-global {
        key 63 {
            secret "ABC123";
            start-time "2011-8-6.06:54:00-0700";
            algorithm hmac-sha-1;
            options isis-enhanced;
        }
    }
    key-chain base-key-inter {
        key 0 {
            secret "$ABC123";
            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**
- [Understanding Hitless Authentication Key Rollover for IS-IS on page 87](#)







## CHAPTER 5

# Configuring Loop-Free Alternate Routes for IS-IS

- [Example: Configuring Link and Node Protection for IS-IS Routes on page 95](#)

### Example: Configuring Link and Node Protection for IS-IS Routes

---

- [Understanding Loop-Free Alternate Routes for IS-IS on page 95](#)
- [Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 99](#)

### 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 12 on page 96](#), 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 12: 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 *Junos OS MPLS Applications Library 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.

**See Also** • [Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 99](#)

### 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 99](#)
- [Overview on page 99](#)
- [Configuration on page 100](#)
- [Verification on page 107](#)

#### 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 13 on page 100 shows the topology used in this example.

Figure 13: 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 100 shows the configuration for all of the devices in Figure 13 on page 100. The section “Step-by-Step Procedure” on page 104 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 <b>[edit]</b> 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 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 </pre> |



```

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-options autonomous-system 65534
set 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 107](#)
- [Checking Which Next-Hop Neighbors Are Designated as Backup Paths to the Destination Node on page 107](#)
- [Checking the Backup Coverage on page 108](#)
- [Checking the Type of Protection Configured on page 109](#)

#### *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    Level1  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    Level1  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 type and 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:
```

```
to0.0
```

```
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.

**See Also** • [Understanding Loop-Free Alternate Routes for IS-IS on page 95](#)

**Related Documentation** • [Example: Configuring BFD for IS-IS on page 77](#)  
• *Understanding BFD for IS-IS*



## CHAPTER 6

# Configuring Segment Routing

- [Understanding Source Packet Routing in Networking \(SPRING\) on page 111](#)
- [Support for SRGB, Anycast, and Prefix Segments in SPRING for IS-IS Protocol on page 114](#)
- [Configuring Anycast and Prefix segments in SPRING for IS-IS Protocol on page 116](#)

## Understanding Source Packet Routing in Networking (SPRING)

Source packet routing or segment routing is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to determine the actual path it should take. In this context, the term 'source' means 'the point at which the explicit route is imposed'. Starting with Junos OS Release 17.2R1, segment routing for IS-IS and OSPFv2 is supported on QFX5100 and QFX10000 switches.

Starting with Junos OS Release 17.3R1, segment routing for IS-IS and OSPFv2 is supported on QFX5110 and QFX5200 switches.

Essentially segment routing engages IGPs like IS-IS and OSPF for advertising two types of network segments or tunnels:

- First, a strict forwarded single-hop tunnel that carries packets over a specific link between two nodes, irrespective of the link cost, referred to as *adjacency segments*.
- Second, a multihop tunnel using shortest path links between two specific nodes, referred to as *node segments*.

Ingress routers can steer a packet through a desired set of nodes and links by pre-appending the packet with an appropriate combination of tunnels.

Segment routing leverages the source routing paradigm. A node steers a packet through an ordered list of instructions, called segments. A segment can represent any instruction, topological or service-based. A segment can have a local semantic to a segment routing node or to a global node within a segment routing domain. Segment routing enforces a flow through any topological path and service chain while maintaining per-flow state only at the ingress node to the segment routing domain. Segment routing can be directly applied to the MPLS architecture with no change on the forwarding plane. A segment is encoded as an MPLS label. An ordered list of segments is encoded as a stack of labels. The segment to process is on the top of the stack. Upon completion of a segment, the



related label is popped from the stack. Segment routing can be applied to the IPv6 architecture, with a new type of routing extension header. A segment is encoded as an IPv6 address. An ordered list of segments is encoded as an ordered list of IPv6 addresses in the routing extension header. The segment to process is indicated by a pointer in the routing extension header. Upon completion of a segment, the pointer is incremented.

Traffic engineering shortcuts are enabled for labeled IS-IS segment routes, when you configure **shortcuts** at the following hierarchy levels:

- **[edit protocols is-is traffic-engineering family inet]** for IPv4 traffic.
- **[edit protocols is-is traffic-engineering family inet6]** for IPv6 traffic.

When source packet routing is deployed in the network, the data center, backbone, and peering devices, switch MPLS packets with a label stack built by the source of the traffic; for example, data center servers. In Junos OS Release 17.4R1, the source-routed traffic co-exists with traffic taking RSVP signaled paths, and source routing is implemented as regular label switching through mpls.0 table using the label operations – pop, swap (to the same label value), and swap-push (for interface protection). In all the cases, traffic can be load balanced between multiple Layer 3 interfaces, or within an aggregate interface. Starting in Junos OS Release 17.4R1, the traffic statistics in a segment routing network can be recorded in an OpenConfig compliant format for the Layer 3 interfaces. The statistics is recorded for the Source Packet Routing in Networking (SPRING) traffic only, excluding RSVP and LDP-signaled traffic, and the family MPLS statistics per interface is accounted for separately. The SR statistics also includes SPRING traffic statistics per link aggregation group (LAG) member, and per segment identifier (SID). To enable recording of segment routing statistics, include *sensor-based-stats* statement at the **[edit protocol isis source-packet-routing]** hierarchy level.

Release History Table

| Release | Description  |
|---------|--|
| 17.4R1  | Starting in Junos OS Release 17.4R1, the traffic statistics in a segment routing network can be recorded in an OpenConfig compliant format for the Layer 3 interfaces. |
| 17.3R1  | Starting with Junos OS Release 17.3R1, segment routing for IS-IS and OSPFv2 is supported on QFX5110 and QFX5200 switches.  |
| 17.2R1  | Starting with Junos OS Release 17.2R1, segment routing for IS-IS and OSPFv2 is supported on QFX5100 and QFX10000 switches.   |

#### Related Documentation

- *OSPF Support for Traffic Engineering*
- *Advertising Label-Switched Paths into OSPFv2*
- *IS-IS Extensions to Support Traffic Engineering*
- *Understanding Forwarding Adjacencies*
- *Understanding LDP-IGP Synchronization*
- *node-segment (Protocols OSPF)*



- *no-advertise-adjacency-segment* (Protocols OSPF)
- *no-source-packet-routing* (Protocols OSPF)
- *sensor-based-stats*
- *sensor* (Junos Telemetry Interface)
- *sensor-based-stats* (Junos Telemetry Interface)
- *show (ospf | ospf3) overview*
- *show (ospf | ospf3) neighbor*
- *show ospf database*
- *show (ospf | ospf3) route*
- *show route table*
- [level \(Global IS-IS\) on page 158](#)
- [show isis database on page 235](#)
- [show isis overview on page 257](#)
- [show isis route on page 261](#)
- [show isis adjacency on page 220](#)
- [source-packet-routing \(Protocols IS-IS and OSPF\) on page 195](#)
- [node-segment \(Protocols IS-IS\) on page 181](#)
- [no-advertise-adjacency-segment \(Protocols IS-IS\) on page 168](#)
- [use-source-packet-routing \(Protocols IS-IS\) on page 209](#)



## Support for SRGB, Anycast, and Prefix Segments in SPRING for IS-IS Protocol

---

Segment routing (SR) or source packet routing in networking (SPRING) is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links without relying on the intermediate nodes in the network to determine the actual path it should take.

Ingress routers can steer a packet through a desired set of nodes and links by prepending the packet with an appropriate combination of tunnels. Segment routing leverages the source routing paradigm. A node steers a packet through an ordered list of instructions, called segments. A segment can have a local semantic to a segment routing node or to a global node within a segment routing domain. Segment routing enforces a flow through any topological path and service chain while maintaining per-flow state only at the ingress node to the segment routing domain.

Essentially segment routing engages interior gateway protocols (IGPs) such as IS-IS and OSPF to advertise two types of network segments or tunnels:

- Adjacency segments - A strict forwarded single-hop tunnel that carries packets over a specific link between two nodes, irrespective of the link cost.
- Prefix segments - A multihop tunnel using equal cost multi-hop aware shortest path links to reach a prefix. Node segment is a special case of prefix segment identifier (SID) that uses shortest path links between two specific nodes.

Segment routing can be directly applied to the MPLS architecture with no change on the forwarding plane. A segment is encoded as an MPLS label. An ordered list of segments is encoded as a stack of labels. Every node in the segment routing domain is allocated labels by the node label manager based on the index range configured for source packet routing. These labels are allocated to the node segment based on the availability of the dynamic label range managed by node label manager. Segment routing global block (SRGB) is the range of label values used in segment routing. Currently, you can configure the index range of the label blocks for segment routing. The default value for the index range is 4096. This causes chunks of 256 label blocks being dynamically allocated by the label manager depending on the availability. Because, there is no predictability of the dynamic label range being allocated to the SRGB, Junos OS allows you to configure SRGB range by configuring the value of the starting label and index range at the **[edit protocols isis source-packet-routing srgb]** hierarchy level. Ensure that the labels in the SRGB label range is not used by any other applications. If a label in the configured label range is used by other application, then a syslog error message **RPD\_ISIS\_SRGBALLOCATIONFAIL** is logged to indicate that the label manager is not able to allocate the requested SRGB label range as a label within the SRGB label range is used by other application. You need to free up the labels that fall under the configured SRGB label range. In order to free up the configured label range, you need to check the label ranges configured at the **[edit protocol mpls label-range]** hierarchy level and re-configure the SRGB label range with label range that is available and restart the routing protocol process (RPD).

Currently, Junos OS allows you to configure SPRING node SID for IPv4 and IPv6 address family per routing instance. This node SID is attached to IPv4 and IPv6 router ID if the



router ID is configured on the loopback interface. Otherwise, the lowest IP address assigned to the loopback interface is chosen as the node SID. You can now configure provide prefix segment identifier (SID) and node SID to prefixes that are advertised in ISIS through policy configuration. Prefix segment index is the index assigned to a specific prefix. This is used by all other remote routers in the network to index the prefix into respective SRGB to derive the segment identifier and to forward the traffic destined for this prefix. The prefix SID supports both IPv4 and IPv6 prefixes. Configuring node SID through policy allows you to choose the loopback address that gets the node SID.



**NOTE:** If the same node SID is configured through policy and through source packet routing at `[edit protocols isis source-packet-routing]` hierarchy level then the SID configured through policy takes precedence.

Once the prefix segment indices are provisioned, the JUNOS device advertises them in one or more of the following ISIS TLV types using a new Prefix-SID Sub-TLV (type 3):

- IP Prefix TLV (type 135)
- MT IP Prefix TLV (type 235)
- IPV6 Prefix Reachability TLV (type 236)
- MT IPV6 Prefix Reachability TLV (type 237)

You can configure the explicit-NULL flag on all the prefixes advertised by a node at the `[edit protocols isis source-packet-routing]` hierarchy level. This enables the P and E bits in all the prefix SID advertisements and allows any upstream neighbor of the node to replace the existing prefix SID with a prefix SID having an explicit null value of 0 for IPv4 and 2 for IPv6 before forwarding the packet.

The IS-IS protocol creates adjacency segments per adjacency, level and, address family (one each for IPv4 and IPv6). An MPLS label is allocated for each adjacency segment that gets created. These labels are allocated once the adjacency status of the segment changes to UP state. In the network, during link flaps a new set of adjacency segment labels are allocated and advertised for the link. In order to ensure that the adjacency segments are retained during adjacency or link flaps for a certain duration, you can configure the hold time of adjacency segment at `[edit protocols isis source-packet-routing]` hierarchy level.

An IGP anycast segment is an IGP prefix segment that identifies a set of routers. An anycast segment or anycast SID enforces forwarding based on the equal-cost multipath-aware shortest-path towards the closest node of the anycast set. Within an anycast group, all the routers advertise the same prefix with the same SID value.

When interior gateway protocols (IGPs) are enabled, the traffic engineering shortcut computations done for normal IPv4 or IPv6 or ISO prefixes gets extended to SPRING node segment tracking routes (in `inet.3` table) and transit routes (in `mpls.0` table). In order to enable IGP traffic engineering shortcuts separately for IS-IS routes and for segment routes per family in MPLS network, you have to configure shortcuts at `[edit`



`protocols isis traffic-engineering family inet-mpls`] hierarchy level for IPv4 traffic and at `[edit protocols isis traffic-engineering family inet6-mpls]` hierarchy level for IPv6 traffic.

- Related Documentation**
- [Configuring Anycast and Prefix segments in SPRING for IS-IS Protocol on page 116](#)
  - [Configuring Segment Routing Global Blocks Label Ranges in SPRING for IS-IS Protocol](#)
  - [Example: Configuring Segment Routing Global Blocks in SPRING for IS-IS to Increase Network Speed](#)
  - [adjacency-segment on page 125](#)
  - [prefix-segment on page 191](#)
  - [source-packet-routing on page 195](#)
  - [srgb on page 200](#)
  - [traffic-engineering on page 205](#)

---

## Configuring Anycast and Prefix segments in SPRING for IS-IS Protocol

Segment routing (SR) or source packet routing in networking (SPRING) is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links without relying on the intermediate nodes in the network to determine the actual path it should take. Segment routing global block (SRGB) is the range of label values used in segment routing. Junos OS allows you to configure prefix segment identifier (SID) and node SID to prefixes that are advertised in IS-IS through policy configuration.

Before you configure SPRING SRGB, prefix SID, and anycast SID for the IS-IS protocol, you must:

- Configure the router interfaces.
- Configure the router ID.
- Configure IS-IS.

To configure device R1 with SPRING SRGB, prefix SID, and anycast SID for IS-IS protocols:

1. Configure the start-label and index-range of SRGB.

```
[edit protocols isis source-packet-routing]
user@host# set srgb start-label start-label-value
user@host# set srgb index-range index-range-value
```

For example, configure SRGB with start-label 800000 and index-range 40000 .

```
[edit protocols isis source-packet-routing]
user@host# set srgb start-label 800000
user@host# set srgb index-range 40000
```

2. Configure the routing policy to match a route (IPv4 or IPv6 ) exactly. Configure the index and the node segment of the prefix segment for a given term and accept the routing policy.



```
[edit policy-options policy-statement policy-name term term-value from]
user@host# set route-filter IP address exact
user@host# set prefix-segment index index-value
user@host# set prefix-segment node-segment
user@host# set accept
```



**NOTE:** Configure node segment as /32 prefix on loopback interface (lo0.0) or on a valid stub interface.

For example, configure the routing policy to match the IPv4 route exactly. Configure the index and the node segment of the prefix segment for a given term and accept the routing policy.

```
[edit policy-options policy-statement policy-name term term-value from]
user@host# set route-filter 198.51.100.1/32 exact
user@host# set prefix-segment index index-value
user@host# set prefix-segment node-segment
user@host# set accept
```

For example, configure the routing policy to match the IPv6 route exactly. Configure the index and the node segment of the prefix segment for a given term and accept the routing policy.

```
[edit policy-options policy-statement policy-name term term-value from]
user@host# set route-filter 2001:db8::/32 exact
user@host# set prefix-segment index index-value
user@host# set prefix-segment node-segment
user@host# set accept
```

3. Configure the index and the node segment of the prefix segment for a given term and accept the routing policy.

```
[edit policy-options policy-statement policy-name term term-value then]
user@host# set prefix-segment index index-value
user@host# set prefix-segment node-segment
user@host# set accept
```

For example, configure the prefix segment with index 1004 and the node segment for term 1 of policy statement prefix SID and accept the routing policy.

```
[edit policy-options policy-statement prefix-sid term 1 then]
user@host# set prefix-segment index 1004
user@host# set prefix-segment node-segment
user@host# set accept
```

4. Configure the routing policy with the same prefix (IPv4 or IPv6 )and same prefix segment on more than one routers for anycast SID.



**NOTE:** For anycast prefix SID, configure prefix SID on loopback interface( lo0.0).

```
[edit policy-options policy-statement prefix-sid term 1 ]
```



```
user@host# set from route-filter IP address exact
user@host# set then prefix-segment index index-value
user@host# set then accept
```

For example, configure IPv4 prefix 198.51.100.1/32 with prefix segment 1000 on two routers R0 and R1 for anycast SID.

```
[edit policy-options policy-statement prefix-sid term 1 ]
user@host# set from route-filter 198.51.100.1/32 exact
user@host# set then prefix-segment index 1000
user@host# set then accept
```

For example, configure IPv6 prefix 2001:db8::/32 with prefix segment 1000 on two routers R0 and R1 for anycast SID.

```
[edit policy-options policy-statement prefix-sid term 1 ]
user@host# set from route-filter 2001:db8::/32 exact
user@host# set then prefix-segment index 2000
user@host# set then accept
```

5. Configure export policy on the IS-IS protocol.

```
[edit protocols isis]
user@host# export prefix-sid
```

6. Configure traffic-engineering shortcuts for IPv4-MPLS family traffic.

```
[edit protocols isis traffic-engineering]
user@host# set family inet-mpls shortcuts
```

7. Configure traffic-engineering shortcuts for IPv6-MPLS family traffic.

```
[edit protocols isis traffic-engineering]
user@host# set family inet6-mpls shortcuts
```

8. Configure explicit NULL to enable E and P bits in all prefix SID advertisements.

```
[edit protocol isis source-packet-routing]
user@host# set explicit-null
```

9. Configure adjacency segment hold time to retain segment adjacency.

```
[edit protocol isis source-packet-routing]
user@host# set adjacency-segment hold-time hold-time
```

For example, configure adjacency segments with 240,000 milliseconds hold time.

```
[edit protocol isis source-packet-routing]
user@host# set adjacency-segment hold-time 240000
```

#### Related Documentation

- [Support for SRGB, Anycast, and Prefix Segments in SPRING for IS-IS Protocol on page 114](#)
- [adjacency-segment on page 125](#)
- [source-packet-routing on page 195](#)



- [traffic-engineering on page 205](#)







## PART 3

# Configuration Statements and Operational Commands

- Configuration Statements on page 123
- Operational Commands on page 211







## CHAPTER 7

# Configuration Statements

- [adjacency-segment on page 125](#)
- [authentication-key \(Protocols IS-IS\) on page 126](#)
- [authentication-key-chain \(Protocols IS-IS\) on page 127](#)
- [authentication-type \(Protocols IS-IS\) on page 128](#)
- [bfd-liveness-detection \(Protocols IS-IS\) on page 129](#)
- [checksum \(Protocols IS-IS\) on page 132](#)
- [csnp-interval on page 133](#)
- [disable \(Protocols IS-IS\) on page 134](#)
- [export on page 135](#)
- [external-preference \(Protocols IS-IS\) on page 136](#)
- [family \(Protocols IS-IS\) on page 137](#)
- [graceful-restart \(Protocols IS-IS\) on page 138](#)
- [hello-authentication-key on page 139](#)
- [hello-authentication-key-chain on page 140](#)
- [hello-authentication-type on page 141](#)
- [hello-interval \(Protocols IS-IS\) on page 142](#)
- [hello-padding on page 144](#)
- [hold-time \(Protocols IS-IS\) on page 146](#)
- [ignore-attached-bit on page 148](#)
- [interface \(Protocols IS-IS\) on page 149](#)
- [ipv4-multicast on page 151](#)
- [ipv4-multicast-metric on page 152](#)
- [ipv6-multicast on page 153](#)
- [ipv6-multicast-metric on page 154](#)
- [ipv6-unicast on page 155](#)
- [ipv6-unicast-metric on page 156](#)
- [isis on page 157](#)
- [level \(Global IS-IS\) on page 158](#)



- [link-protection \(Protocols IS-IS\) on page 159](#)
- [loose-authentication-check on page 160](#)
- [lsp-interval on page 161](#)
- [lsp-lifetime on page 163](#)
- [max-areas on page 164](#)
- [mesh-group \(Protocols IS-IS\) on page 165](#)
- [metric \(Protocols IS-IS\) on page 166](#)
- [no-adjacency-holddown on page 167](#)
- [no-advertise-adjacency-segment \(Protocols IS-IS\) on page 168](#)
- [no-authentication-check on page 169](#)
- [no-csnp-authentication on page 170](#)
- [no-eligible-backup \(Protocols IS-IS\) on page 171](#)
- [no-hello-authentication on page 172](#)
- [no-ipv4-multicast on page 173](#)
- [no-ipv4-routing on page 174](#)
- [no-ipv6-multicast on page 175](#)
- [no-ipv6-routing on page 176](#)
- [no-ipv6-unicast on page 177](#)
- [no-psnp-authentication on page 178](#)
- [no-unicast-topology on page 179](#)
- [node-link-protection \(Protocols IS-IS\) on page 180](#)
- [node-segment \(Protocols IS-IS\) on page 181](#)
- [overload \(Protocols IS-IS\) on page 183](#)
- [passive \(Protocols IS-IS\) on page 186](#)
- [point-to-point on page 188](#)
- [preference \(Protocols IS-IS\) on page 189](#)
- [prefix-export-limit \(Protocols IS-IS\) on page 190](#)
- [prefix-segment on page 191](#)
- [priority \(Protocols IS-IS\) on page 192](#)
- [reference-bandwidth \(Protocols IS-IS\) on page 193](#)
- [rib-group \(Protocols IS-IS\) on page 194](#)
- [source-packet-routing \(Protocols IS-IS and OSPF\) on page 195](#)
- [source-packet-routing-disable \(Protocols IS-IS\) on page 197](#)
- [spf-options \(Protocols IS-IS\) on page 198](#)
- [srgb on page 200](#)
- [topologies \(Protocols IS-IS\) on page 201](#)
- [traceoptions \(Protocols IS-IS\) on page 202](#)




- [traffic-engineering \(Protocols IS-IS\) on page 205](#)
- [use-source-packet-routing \(Protocols IS-IS\) on page 209](#)
- [wide-metrics-only on page 210](#)

## adjacency-segment

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | adjacency-segment {<br>hold-time <i>hold-time</i> ;<br>}  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>name</i> protocols isis <a href="#">source-packet-routing</a> ],<br>[edit logical-systems <i>name</i> routing-instances <i>name</i> protocols isis <a href="#">source-packet-routing</a> ],<br>[edit protocols isis source-packet-routing],<br>[edit routing-instances <i>name</i> protocols isis <a href="#">source-packet-routing</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 17.2R1 for MX Series routers, PTX Series routers, QFX5100 switches, and QFX10000 switches.<br>Statement introduced in Junos OS Release 17.3R1 for QFX5110 and QFX5200 switches.  |
| <b>Description</b>              | Configure attributes for adjacency segments in source packet routing in networking (SPRING), or configure segment routing (SR) to ensure that the adjacency segment identifiers are retained during adjacency or link flaps. The adjacency segments are not released immediately and are retained for the configured hold time duration.  |
| <b>Options</b>                  | <b>hold-time</b> —Duration, in milliseconds, to retain adjacency segments after isolating from an interface.<br><b>Range:</b> 180,000 through 900,000   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Support for SRGB, Anycast, and Prefix Segments in SPRING for IS-IS Protocol on page 114</a></li> <li>• <a href="#">Configuring Anycast and Prefix segments in SPRING for IS-IS Protocol on page 116</a></li> <li>• <a href="#">source-packet-routing on page 195</a></li> <li>• <a href="#">traffic-engineering on page 205</a></li> <li>• <a href="#">prefix-segment on page 191</a></li> </ul> |



## authentication-key (Protocols IS-IS)

|   |  |
|---|--|
| <b>Syntax</b>   | authentication-key <i>key</i> ;  |
| <b>Hierarchy Level</b>  | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>level</b> <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>level</b> <i>level-number</i>],</p> <p>[edit protocols isis <b>level</b> <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>level</b> <i>level-number</i>]</p>   |
| <b>Release Information</b>  | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>  |
| <b>Description</b>  | <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 <b>authentication-type</b> 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> |
| <b>Default</b>  | If you do not include this statement and the <b>authentication-type</b> statement, IS-IS authentication is disabled.   |
| <b>Options</b>  | <p><b>key</b>—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 (" ").</p>   |
| <div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p><b>CAUTION:</b> A simple password for authentication is truncated if it exceeds 254 characters.</p> </div> </div> |  |
| <b>Required Privilege Level</b>   | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Hitless Authentication Key Rollover for IS-IS on page 87</a></li> <li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 88</a></li> </ul>  |



## authentication-key-chain (Protocols IS-IS)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | authentication-key-chain <i>key-chain-name</i> ;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>name</i> protocols isis level <i>level-number</i> ],<br>[edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols isis level <i>level-number</i> ],<br>[edit protocols isis level <i>level-number</i> ],<br>[edit routing-instances <i>instance-name</i> protocols isis level <i>level-number</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.2.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | Apply and enable an authentication keychain to the routing device.  |
| <b>Options</b>                  | <b>key-chain</b> —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 (" ").   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 88</a></li> <li>• <a href="#">Example: Configuring Router Authentication for BGP</a></li> <li>• <a href="#">Example: Configuring BFD Authentication for Securing Static Routes</a></li> <li>• <a href="#">Configuring Authentication Key Updates</a></li> <li>• <a href="#">Understanding Hitless Authentication Key Rollover for IS-IS on page 87</a></li> </ul> |



## authentication-type (Protocols IS-IS)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>authentication-type <i>authentication</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | Enable authentication and specify the authentication scheme for IS-IS. If you enable authentication, you must specify a password by including the <b>authentication-key</b> statement.   |
| <b>Default</b>                  | If you do not include this statement and the <b>authentication-key</b> statement, IS-IS authentication is disabled.  |
| <b>Options</b>                  | <b><i>authentication</i></b> —Authentication scheme: <ul style="list-style-type: none"><li>• <b>md5</b>—Use HMAC authentication in combination with MD5. HMAC-MD5 authentication is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>.</li><li>• <b>simple</b>—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.</li></ul> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">authentication-key on page 126</a></li><li>• <a href="#">no-authentication-check on page 169</a></li><li>• <a href="#">Understanding Hitless Authentication Key Rollover for IS-IS on page 87</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 88</a></li></ul>  |



## bfd-liveness-detection (Protocols IS-IS)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <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>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a>],</p> <p>[edit protocols isis <a href="#">interface interface-name</a>],</p> <p>[edit protocols isis <a href="#">interface interface-name</a> family inet inet6],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a>]</p>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>detection-time threshold</b> and <b>transmit-interval threshold</b> options added in Junos OS Release 8.2.</p> <p>Support for logical systems introduced in Junos OS Release 8.3.</p> <p><b>no-adaptation</b> statement introduced in Junos OS Release 9.0.</p> <p><b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> options introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| <b>Description</b>         | Configure bidirectional failure detection timers and authentication.   |



**NOTE:** Starting with Junos OS Release 16.1, you can configure IS-IS BFD sessions for the IPv6 address family. Therefore, this configuration statement is also available under the [edit protocols isis [interface interface-name](#) family inet|inet6] hierarchy level.



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 does not affect traffic flow on the routing device.



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

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

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

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

**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

The remaining statements are explained separately. See [CLI Explorer](#).

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

**Related Documentation**

- [Example: Configuring BFD for IS-IS on page 77](#)
- [Understanding BFD Authentication for IS-IS on page 75](#)

---

## checksum (Protocols IS-IS)

---

**Syntax** checksum;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols isis [interface](#) *interface-name*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis [interface](#) *interface-name*],  
[edit protocols isis [interface](#) *interface-name*],  
[edit routing-instances *routing-instance-name* 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, *Optional Checksums in Intermediate System to Intermediate System (ISIS)*.

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.



## csnp-interval

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>csnp-interval (seconds   disable);</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>  |
| <b>Description</b>              | <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.</p>  |
|                                 | <p> <b>NOTE:</b> The <code>csnp-interval</code> statement is effective only when configured on LAN interfaces.</p>  |
|                                 | <p>To configure the interface not to send any CSNPs, specify the <b>disable</b> option.</p>  |
| <b>Default</b>                  | 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.   |
| <b>Options</b>                  | <p><b>disable</b>—Do not send CSNPs on this interface.</p> <p><i>seconds</i>—Number of seconds between the sending of CSNPs.</p> <p><b>Range:</b> 1 through 65,535 seconds</p> <p><b>Default:</b> 10 seconds on LAN broadcast links</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces</i></li> <li>• <i>Understanding the Transmission Frequency for CSNPs on IS-IS Interfaces</i></li> </ul>   |




## disable (Protocols IS-IS)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | disable;   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>isis</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <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 <b>traffic-engineering</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>isis</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <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 <b>traffic-engineering</b>],</p> <p>[edit protocols <b>isis</b>],</p> <p>[edit protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis <b>traffic-engineering</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>isis</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <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 <b>traffic-engineering</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>  |
| <b>Description</b>              | <p>Disable IS-IS on the routing device, on an interface, or on a level.</p> <p>At the [edit protocols isis <b>traffic-engineering</b>] hierarchy level, disable IS-IS support for traffic engineering.</p> <p>Enabling IS-IS on an interface (by including the <b>interface</b> 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 <b>disable</b> statement), and not actually having IS-IS run on an interface (by including the <b>passive</b> statement) are mutually exclusive states.</p>  |
| <b>Default</b>                  | <p>IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which <b>family iso</b> is enabled.</p> <p>IS-IS support for traffic engineering is enabled.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |



- Related Documentation**
- [Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 17](#)
  - [IS-IS Overview on page 3](#)

## export

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| <b>Description</b>              | <p>Apply one or more policies to routes being exported from the routing table into IS-IS.</p> <p>All routing protocols store the routes that they learn in the routing table. The routing table uses this collected route information to determine the active routes to destinations. The routing table then installs the active routes into its forwarding table and exports them into the routing protocols. It is these exported routes that the protocols advertise.</p> <p>For each protocol, you control which routes the protocol stores in the routing table and which routes the routing table exports into the protocol from the routing table by defining a <i>routing policy</i> for that protocol.</p> |
|                                 | <p> <b>NOTE:</b> For IS-IS, you cannot apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol can easily lead to an inconsistent topology database.</p>  |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    |   |



## external-preference (Protocols IS-IS)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>external-preference <i>preference</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| <b>Description</b>              | Configure the preference of external routes.  |
| <b>Options</b>                  | <i>preference</i> —Preference value.<br><b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ )<br><b>Default:</b> 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Route Preferences Overview</i></li> <li>• <a href="#">Example: Redistributing OSPF Routes into IS-IS on page 28</a></li> <li>• <i>Example: Configuring a Routing Policy to Redistribute BGP Routes with a Specific Community Tag into IS-IS</i></li> <li>• <i>Understanding BGP Communities, Extended Communities, and Large Communities as Routing Policy Match Conditions</i></li> <li>• <i>Understanding Routing Policies</i></li> <li>• <a href="#">preference on page 189</a></li> </ul> |



## family (Protocols IS-IS)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> family inet {     shortcuts {         multicast-rpf-routes;     } } family inet-mpls {     shortcuts; } family inet6 {     shortcuts; } family inet6-mpls {     shortcuts; } </pre>   |
| <b>Hierarchy Level</b>          | <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>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.3.</p> <p>Support for IPv6 for IGP shortcuts introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Support for <b>inet-mpls</b> and <b>inet6-mpls</b> statements introduced in Junos OS Release 17.2 for the MX Series routers, PTX Series routers, QFX5100, and QFX10000 switches.</p> |
| <b>Description</b>              | Configure the address family for traffic engineering IS-IS interior gateway protocol (IGP) shortcuts.   |
| <b>Options</b>                  | <p>inet—IPv4 address family</p> <p>inet-mpls—IPv4-MPLS address family</p> <p>inet6—IPv6 address family</p> <p>inet6-mpls—IPv6-MPLS address family</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |



## graceful-restart (Protocols IS-IS)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>graceful-restart {<br/>    disable;<br/>    helper-disable;<br/>    restart-duration <i>seconds</i>;<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>              | <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 <a href="#">[edit routing-options]</a> 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"><li>• The ingress router of the label-switched path (LSP), which passes through the DR, to break the LSP.</li><li>• The ingress router to re-signal the LSP.</li></ul> |
| <b>Options</b>                  | <p><b>disable</b>—Disable graceful restart for IS-IS.</p> <p><b>helper-disable</b>—Disable graceful restart helper capability. Helper mode is enabled by default.</p> <p><b>restart-duration <i>seconds</i></b>—Time period for the restart to last, in seconds.<br/><b>Range:</b> 30 through 300 seconds<br/><b>Default:</b> 30 seconds</p>   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Routing Protocols Graceful Restart</i></li></ul>  |



## hello-authentication-key

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>hello-authentication-key password;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>  |
| <b>Description</b>              | 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 <b>hello-authentication-type</b> statement.  |
| <b>Default</b>                  | 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.  |
| <b>Options</b>                  | <p><b>password</b>—Authentication password. The password can be up to 255 characters.</p> <p>Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">authentication-key on page 126</a></li> <li>• <a href="#">authentication-type on page 128</a></li> <li>• <a href="#">hello-authentication-type on page 141</a></li> </ul>   |




## hello-authentication-key-chain

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | hello-authentication-key-chain <i>key-chain-name</i> ;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit routing-instances <i>instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.2.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | Apply an authentication keychain to the IS-IS interface.  |
| <b>Options</b>                  | <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 (" ").  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Hitless Authentication Key Rollover for IS-IS on page 87</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 88</a></li></ul>  |



## hello-authentication-type

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | hello-authentication-type (md5   simple);  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>  |
| <b>Description</b>              | <p>Enable authentication on an interface for hello packets. If you enable authentication on hello packets, you must specify a password by including the <b>hello-authentication-key</b> statement.</p> <p>You can configure authentication for a given IS-IS level on an interface. On a point-to-point link, if you enable hello authentication for both IS-IS levels, the password configured for Level 1 is used for both levels.</p>   |
|                                 | <div>  <p><b>CAUTION:</b> 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>  |
| <b>Default</b>                  | 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.  |
| <b>Options</b>                  | <p><b>md5</b>—Specifies Message Digest 5 as the packet verification type.</p> <p><b>simple</b>—Specifies simple authentication as the packet verification type.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">authentication-key on page 126</a></li> <li>• <a href="#">authentication-type on page 128</a></li> <li>• <a href="#">hello-authentication-key on page 139</a></li> </ul>  |



## hello-interval (Protocols IS-IS)

**Syntax** `hello-interval seconds;`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols isis interface *interface-name* level *level-number*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*],  
[edit protocols isis interface *interface-name* level *level-number*],  
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* 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.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Modify the frequency with which the routing device sends hello packets out of an interface, in seconds.

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.

You can send out hello packets in subsecond intervals. To send out hello packets every 333 milliseconds, set the **hold-time** value to 1.

**Options** *seconds*—Frequency of transmission for hello packets.

**Range:** 1 through 20,000 seconds

**Default:** 3 seconds (for designated intermediate system [DIS] routers), 9 seconds (for non-DIS routers)



**NOTE:** When elected as a designated intermediate system [DIS] router on any LAN adjacency, the hello and hold-timer intervals are scaled down by a factor of 3. This means the default values of 9 and 27 seconds of ISIS hello and hold time intervals are scaled down to 3 and 9 seconds for LAN adjacencies. During switchovers, this hello interval is too short to form LAN adjacencies. Therefore, you can configure one of the following solutions:

- Set the hello and hold time interval for LAN adjacencies to 30 seconds and 90 seconds respectively on both the DIS and a neighboring router.
- Convert the LAN interfaces to point-to-point IS-IS interfaces.





**BEST PRACTICE:** On QFX10000 switches, we strongly recommend that you configure all IS-IS interfaces, including peer interfaces, as point-to-point interfaces. If you do not, you might experience session flaps, that is, IS-IS sessions that go down and then come back up, when IS-IS is configured in virtual routing and forwarding (VRF) instances. When you scale IS-IS in any scenario, you might also experience scaling issues if you do not configure IS-IS interfaces as point-to-point interfaces.

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

**Related Documentation** • *hold-time*



## hello-padding

---

**Syntax** hello-padding (adaptive | loose | strict);

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

**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.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure padding on hello packets to accommodate asymmetrical maximum transfer units (MTUs) from different hosts.

This helps to prevent a premature adjacency Up state when one routing device's MTU does not meet the requirements to establish the adjacency.

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.

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.

There are four types of hello padding:

- 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**    routing—To view this statement in the configuration.

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



## hold-time (Protocols IS-IS)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>hold-time seconds;</code>   |
| <b>Hierarchy Level</b>     | <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>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>   |
| <b>Description</b>         | <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> <p>For systems configured with graceful routing switchover (GRES) with Graceful Restart, the hold time for Master and Backup Routing Engines should be set to a value higher than 40 seconds. This ensures that adjacencies between the Routing Engine and the neighboring peer 'helper' routers do not time out, stopping graceful restart, and all traffic.</p> |
| <b>Options</b>             | <p><b>seconds</b>—Hold-time value, in seconds.</p> <p><b>Range:</b> 3 through 65,535 seconds, or 1 to send out hello packets every 333 milliseconds</p> <p><b>Default:</b> 9 seconds (for designated intermediate system [DIS] routers), 27 seconds (for non-DIS routers; three times the default hello interval)</p>   |



**NOTE:** When elected as a designated intermediate system [DIS] router on any LAN adjacency, the hello and hold-timer intervals are scaled down by a factor of 3. This means the default values of 9 and 27 seconds of ISIS hello and hold time intervals are scaled down to 3 and 9 seconds for LAN adjacencies. During switchovers, this hold time is too short to form LAN adjacencies. Therefore, you can configure one of the following solutions:

- Set the hello and hold time interval for LAN adjacencies to 30 seconds and 90 seconds respectively on both the DIS router and a neighboring router.
- Convert the LAN interfaces to point-to-point IS-IS interfaces.





**BEST PRACTICE:** On QFX10000 switches, we strongly recommend that you configure all IS-IS interfaces, including peer interfaces, as point-to-point interfaces. If you do not, you might experience session flaps, that is, IS-IS sessions that go down and then come back up, when IS-IS is configured in virtual routing and forwarding (VRF) instances. When you scale IS-IS in any scenario, you might also experience scaling issues if you do not configure IS-IS interfaces as point-to-point interfaces.

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Graceful Routing Engine Switchover</i></li><li>• <a href="#">Example: Configuring IS-IS on page 11</a></li><li>• <i>Example: Configuring IS-IS for GRES with Graceful Restart</i></li><li>• <i>Understanding IS-IS Configuration</i></li><li>• <a href="#">hello-interval on page 142</a></li></ul> |



## ignore-attached-bit

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | ignore-attached-bit;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <i>isis</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i> ],<br>[edit protocols <i>isis</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | <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 <b>ignore-attached-bit</b> 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> |
| <b>Default</b>                  | The <b>ignore-attached-bit</b> statement is disabled by default.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |



## 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;
    flood-group;
    csnp-interval (seconds | disable);
    hello-padding (adaptive | loose | strict);
    interface-group-holddown-delay
    ldp-synchronization {
        disable;
        hold-time seconds;
    }
    lsp-interval milliseconds;
    max-hello-size size;
    mesh-group (value | blocked);
    no-adjacency-holddown;
    no-eligible-remote-backup;
    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;
    }
}

```

|                                 |   |
|---------------------------------|---|
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>],</p> <p>[edit protocols <a href="#">isis</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p><b>no-eligible-remote-backup</b> option introduced in Junos OS Release 14.2 for the MX Series.</p> <p><b>interface-group-holddown-delay</b> option introduced in Junos OS Release 15.2 for the MX Series.</p> <p><b>max-hello-size</b> option introduced in Junos OS Release 17.1 for the MX Series.</p> |
| <b>Description</b>              | <p>Configure interface-specific IS-IS properties. To configure more than one interface, include the <b>interface</b> statement multiple times.</p> <p>Enabling IS-IS on an interface (by including the <b>interface</b> statement at the [edit protocols <a href="#">isis</a>] or the [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>] hierarchy level), disabling it (by including the <b>disable</b> statement), and not actually having IS-IS run on an interface (by including the <b>passive</b> statement) are mutually exclusive states.</p>                   |
| <b>Options</b>                  | <p><b>all</b>—Have Junos OS create IS-IS interfaces automatically. If you include this option, disable IS-IS on the management interface (fxp0).</p> <p><b>interface-name</b>—Name of an interface. Specify the full interface name, including the physical and logical address components.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding IS-IS Configuration</a></li> <li>• <a href="#">Example: Configuring IS-IS on page 11</a></li> <li>• <a href="#">Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups</a></li> <li>• <a href="#">Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 17</a></li> </ul>  |



## ipv4-multicast

|  |   |
|--|---|
| <b>Syntax</b>  | ipv4-multicast;   |
| <b>Hierarchy Level</b>   | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit protocols isis <a href="#">topologies</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ] |
| <b>Release Information</b>   | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| <b>Description</b>   | Configure alternate IPv4 multicast topologies.  |
| <div>  <p><b>NOTE:</b> 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> |   |
| <b>Default</b>   | Multicast topologies are disabled.  |
| <b>Required Privilege Level</b>  | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 47</a></li> <li>• <a href="#">IS-IS Multicast Topologies Overview on page 46</a></li> </ul>  |



## ipv4-multicast-metric

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>ipv4-multicast-metric <i>metric</i>;</code>   |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code><br><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><br><code>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| <b>Description</b>              | Specify the multicast topology metric value for the level.  |
| <b>Options</b>                  | <b><i>metric</i></b> —Metric value.<br><b>Range:</b> 0 through 16,777,215   |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 47</a></li><li>• <a href="#">IS-IS Multicast Topologies Overview on page 46</a></li></ul>   |



## ipv6-multicast

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | ipv6-multicast;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit protocols isis <a href="#">topologies</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | Configure alternate IPv6 multicast topologies.  |
| <b>Default</b>                  | Multicast topologies are disabled.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 47</a></li> <li>• <a href="#">IS-IS Multicast Topologies Overview on page 46</a></li> </ul>  |




## ipv6-multicast-metric

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>ipv6-multicast-metric <i>metric</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[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> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | Specify the IPv6 alternate multicast topology metric value for the level.   |
| <b>Options</b>                  | <i>metric</i> —Metric value.<br><b>Range:</b> 0 through 16,777,215  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 47</a></li><li>• <a href="#">IS-IS Multicast Topologies Overview on page 46</a></li></ul>   |



## ipv6-unicast

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | ipv6-unicast;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit protocols isis <a href="#">topologies</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>              | <p>Configure alternate IPv6 unicast topologies.</p> <p>This statement causes IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to inet6.0.</p> <hr/> <div>  <p><b>NOTE:</b> 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> <hr/> |
| <b>Default</b>                  | IPv6 unicast topologies are disabled.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li> <li>• <a href="#">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li> </ul>  |



## ipv6-unicast-metric

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>ipv6-unicast-metric <i>metric</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[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> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | 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.  |
| <b>Options</b>                  | <i>metric</i> —Metric value.<br><b>Range:</b> 0 through 16,777,215  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li><li>• <a href="#">Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li></ul>  |



## isis

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | isis { ... }   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols]        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| <b>Description</b>              | Enable IS-IS routing on the routing device or for a routing instance.<br><br>The <b>isis</b> statement is the one statement you must include in the configuration to run IS-IS on the routing device or in a routing instance.   |
| <b>Default</b>                  | IS-IS is disabled on the routing device.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS on page 11</a></li> <li>• <a href="#">Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 17</a></li> </ul>  |



## level (Global IS-IS)

---

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre>level <i>level-number</i> {<br/>    authentication-key <i>key</i>;<br/>    authentication-key-chain (Protocols IS-IS) <i>key-chain-name</i>;<br/>    authentication-type <i>type</i>;<br/>    disable;<br/>    external-preference <i>preference</i>;<br/>    no-csnp-authentication;<br/>    no-hello-authentication;<br/>    no-psnp-authentication;<br/>    preference <i>preference</i>;<br/>    purge-originator {<br/>        empty self;<br/>    }<br/>    wide-metrics-only;<br/>}</pre>   |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]   |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| <b>Description</b>         | <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> |
| <b>Options</b>             | <p><i>level-number</i>—IS-IS level number.</p> <p><b>Values:</b> 1 or 2</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>   |



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

**Related Documentation**

- [Understanding IS-IS Configuration](#)
- [Example: Configuring IS-IS on page 11](#)
- [Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups](#)
- [Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 17](#)

## link-protection (Protocols IS-IS)

**Syntax** link-protection;

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

**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.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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**

- [Understanding Loop-Free Alternate Routes for IS-IS on page 95](#)
- [Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 99](#)
- [node-link-protection on page 180](#)



## loose-authentication-check

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | loose-authentication-check;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| <b>Description</b>              | Allow the use of MD5 authentication without requiring network-wide deployment.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Hitless Authentication Key Rollover for IS-IS on page 87</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 88</a></li></ul>  |



## **lsp-interval**

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>lsp-interval <i>milliseconds</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],<br/>         [edit protocols isis <b>interface</b> <i>interface-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/>         Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>         Statement introduced in Junos OS Release 12.1 for the QFX Series.<br/>         Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>   |
| <b>Description</b>              | <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 <b>lsp-interval</b> 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> |
| <b>Default</b>                  | By default, the routing device sends one link-state PDU out an interface every 100 milliseconds.  |
| <b>Options</b>                  | <p><b><i>milliseconds</i></b>—Number of milliseconds between the sending of link-state PDUs. Specifying a value of 0 blocks all link-state PDU transmission.</p> <p><b>Range:</b> 0 through 1000 milliseconds</p> <p><b>Default:</b> 100 milliseconds</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i></li> <li>• <i>Understanding Link-State PDU Throttling for IS-IS Interfaces</i></li> <li>• <i>Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces</i></li> </ul>  |



- *Understanding the Transmission Frequency for CSNPs on IS-IS Interfaces*



## **lsp-lifetime**

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>lsp-lifetime <i>seconds</i>;</code>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],<br/> [edit protocols <i>isis</i>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/> Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/> Statement introduced in Junos OS Release 12.1 for the QFX Series.<br/> Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>   |
| <b>Description</b>              | <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 <a href="#">show isis overview</a> command displays the link-state PDU lifetime.</p> |
| <b>Default</b>                  | 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.   |
| <b>Options</b>                  | <p><b><i>seconds</i></b>—link-state PDU lifetime, in seconds.<br/> <b>Range:</b> 350 through 65,535 seconds<br/> <b>Default:</b> 1200 seconds</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/> routing-control—To add this statement to the configuration.</p>  |



- Related Documentation**
- *Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces*
  - *Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces*
  - *Understanding Link-State PDU Throttling for IS-IS Interfaces*
  - *Understanding the Transmission Frequency for CSNPs on IS-IS Interfaces*
  - [https://www.juniper.net/us/en/training/certification/JNCIP\\_studyguide.pdf](https://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf)

---

## max-areas

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>max-areas <i>number</i>;</code>  |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</code><br><code>[edit protocols <i>isis</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</code>   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | <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> |
| <b>Options</b>                  | <p><b><i>number</i></b>—Maximum number of areas to include in the IS-IS hello (IIH) PDUs and link-state PDUs.</p> <p><b>Range:</b> 3 through 36</p> <p><b>Default:</b> 3</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups</i></li><li>• <i>Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 17</i></li></ul>  |



## mesh-group (Protocols IS-IS)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>mesh-group (blocked   <i>value</i>);</code>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>   |
| <b>Description</b>              | <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> |
| <b>Options</b>                  | <p><b>blocked</b>—Configure the interface so that it does not flood link-state PDUs.</p> <p><b>value</b>—Number that identifies the mesh group.</p> <p><b>Range:</b> 1 through 4,294,967,295 (<math>2^{32} - 1</math>; 32 bits are allocated to identify a mesh group)</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Mesh Groups of IS-IS Interfaces</i></li> <li>• <i>Understanding IS-IS Mesh Groups</i></li> </ul>   |



## metric (Protocols IS-IS)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>metric <i>metric</i>;</code>   |
| <b>Hierarchy Level</b>     | <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>   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>  |
| <b>Description</b>         | <p>Specify the IPv4 unicast topology metric value for the level. The IS-IS interface metrics for the IPv6 topology can be configured independently of the IPv4 metrics.</p> <p>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 (<math>2^{24} - 1</math>) if you are using wide metrics.</p> <p>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.</p> |

Table 5 on page 166 depicts IS-IS route export metric default values.

**Table 5: 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.



|                                 |   |
|---------------------------------|---|
| <b>Options</b>                  | <p><i>metric</i>—Metric value.</p> <p><b>Range:</b> 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)</p> <p><b>Default:</b> 10 (for all interfaces except lo0), 0 (for the lo0 interface)</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i></li> <li>• <i>Understanding Wide IS-IS Metrics for Traffic Engineering</i></li> <li>• <i>te-metric</i></li> <li>• <a href="#">wide-metrics-only on page 210</a></li> </ul> |

## no-adjacency-holddown

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-adjacency-holddown;  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],</p> <p>[edit protocols <i>isis</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>   |
| <b>Description</b>              | <p>Disable the hold-down timer for IS-IS adjacencies.</p> <p>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.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">hold-time on page 146</a></li> </ul>   |



## no-advertise-adjacency-segment (Protocols IS-IS)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-advertise-adjacency-segment;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ], |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 16.2 for MX Series routers and PTX Series routers.<br>Statement introduced in Junos OS Release 17.2R1 for QFX5100 and QFX10000 switches.<br>Statement introduced in Junos OS Release 17.3R1 for QFX5100 and QFX5200 switches.   |
| <b>Description</b>              | Disable advertising of the adjacency segment on all levels for the specified interface.  |
| <b>Default</b>                  | Enabled  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">isis on page 157</a></li><li>• <a href="#">source-packet-routing on page 195</a></li></ul>   |



## no-authentication-check

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-authentication-check;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| <b>Description</b>              | Generate authenticated packets and check the authentication on received packets, but do not reject packets that cannot be authenticated.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">csnp-interval on page 133</a></li> <li>• <a href="#">hello-authentication-type on page 141</a></li> </ul>  |



## no-csnp-authentication

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-csnp-authentication;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | Suppress authentication check on complete sequence number PDU (CSNP) packets.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">csnp-interval on page 133</a></li></ul>  |



## no-eligible-backup (Protocols IS-IS)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-eligible-backup;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis interface <i>interface-name</i> ],<br>[edit protocols isis interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | Exclude the specified interface as a backup interface for IS-IS interfaces on which link protection or node-link protection is enabled.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Loop-Free Alternate Routes for IS-IS on page 95</a></li> <li>• <a href="#">Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 99</a></li> <li>• <a href="#">link-protection on page 159</a></li> <li>• <a href="#">node-link-protection on page 180</a></li> </ul>   |



## no-hello-authentication

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-hello-authentication;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | Suppress authentication check on complete sequence number hello packets.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">hello-authentication-type on page 141</a></li></ul>  |




## no-ipv4-multicast

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-ipv4-multicast;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | Exclude an interface from IPv4 multicast topologies.   |
| <b>Default</b>                  | Multicast topologies are disabled.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 47</a></li> <li>• <a href="#">IS-IS Multicast Topologies Overview on page 46</a></li> </ul>   |



## no-ipv4-routing

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-ipv4-routing;  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>],<br/> [edit protocols <a href="#">isis</a>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/> Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/> Statement introduced in Junos OS Release 12.1 for the QFX Series.<br/> Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>   |
| <b>Description</b>              | <p>Disable IP version 4 (IPv4) routing.</p> <p>Disabling IPv4 routing has the following results:</p> <ul style="list-style-type: none"> <li>• The routing device does not advertise the network layer protocol identifier (NLPID) for IPv4 in the Junos OS link-state PDU fragment zero.</li> <li>• The routing device does not advertise any IPv4 prefixes in Junos OS link-state PDUs.</li> <li>• The routing device does not advertise the NLPID for IPv4 in Junos OS hello packets.</li> <li>• The routing device does not advertise any IPv4 addresses in Junos OS hello packets.</li> <li>• The routing device does not calculate any IPv4 routes.</li> </ul> |
|                                 | <p> <b>NOTE:</b> 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>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/> routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li> <li>• <a href="#">Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li> </ul>   |



---

## no-ipv6-multicast

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-ipv6-multicast;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>              | Exclude an interface from the IPv6 multicast topologies.   |
| <b>Default</b>                  | Multicast topologies are disabled.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 47</a></li><li>• <a href="#">IS-IS Multicast Topologies Overview on page 46</a></li></ul>  |



## no-ipv6-routing

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-ipv6-routing;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | <p>Disable IP version 6 (IPv6) routing.</p> <p>Disabling IPv6 routing has the following results:</p> <ul style="list-style-type: none"><li>• The routing device does not advertise the network layer protocol identifier (NLPID) for IPv6 in the Junos OS link-state PDU fragment zero.</li><li>• The routing device does not advertise any IPv6 prefixes in Junos OS link-state PDUs.</li><li>• The routing device does not advertise the NLPID for IPv6 in Junos OS hello packets.</li><li>• The routing device does not advertise any IPv6 addresses in Junos OS hello packets.</li><li>• The routing device does not calculate any IPv6 routes.</li></ul> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li><li>• <a href="#">Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li></ul>  |



## no-ipv6-unicast

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-ipv6-unicast;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>              | Exclude an interface from the IPv6 unicast topologies. This enables you to exercise control over the paths that unicast data takes through a network.  |
| <b>Default</b>                  | IPv6 unicast topologies are disabled.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li> <li>• <a href="#">Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li> </ul>  |



## no-psnp-authentication

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-psnp-authentication;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | Suppress authentication check on partial sequence number PDU (PSNP) packets.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring IS-IS Authentication on page 26</a></li></ul>  |



## no-unicast-topology

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-unicast-topology;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | Exclude an interface from the IPv4 unicast topologies.   |
| <b>Default</b>                  | IPv4 unicast topologies are disabled.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 47</a></li> <li>• <a href="#">IS-IS Multicast Topologies Overview on page 46</a></li> </ul>   |



## node-link-protection (Protocols IS-IS)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | node-link-protection;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis interface <i>interface-name</i> ],<br>[edit protocols isis interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | 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.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Loop-Free Alternate Routes for IS-IS on page 95</a></li><li>• <a href="#">Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 99</a></li><li>• <a href="#">link-protection on page 159</a></li></ul>   |



## node-segment (Protocols IS-IS)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre>node-segment {     index-range <i>index range</i>;     ipv4-index <i>index</i>;     ipv6-index <i>index</i>; }</pre>   |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis source-packet-routing</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis source-packet-routing</a> ],<br>[edit protocols <a href="#">isis source-packet-routing</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis source-packet-routing</a> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 16.2R1 for MX Series routers.<br>Statement introduced in Junos OS Release 16.1X65 for the PTX1000 routers.<br>Statement introduced in Junos OS Release 17.2R1 for QFX5100 and QFX10000 switches.<br>Statement introduced in Junos OS Release 17.3R1 for QFX5110 and QFX5200 switches.  |
| <b>Description</b>         | Enable source packet routing in networking (SPRING) at all levels. SPRING or segment routing is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to determine the actual path it should take.  |



**NOTE:** Provisioning the IPv4 and IPv6 node segment index is allowed per routing-instance, and will NOT be allowed per IS-IS level. Node segment index is attached to the IPv4 and IPv6 router-id, if the router-ids are configured on the loopback interface. Else, lowest IP address on the loopback is chosen to attach the node-sid.

**Options** **index-range *index range***— Range of node segment indices allowed.

**Default:** 4096

**Range:** 32 through 16384

**ipv4-index *index***— IPv4 node segment index.

**Range:** 0 through 199999



**NOTE:** Starting with Junos OS Release 17.2, the maximum index for IPv4 node segment index is 199999.

**ipv6-index *index***— IPv6 node segment index.

**Range:** 0 through 199999





NOTE: Starting with Junos OS Release 17.2, the maximum index for IPv6 node segment index is 199999.

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

**Related Documentation** • [isis on page 157](#)



## overload (Protocols IS-IS)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre> overload {     advertise-high-metrics;     allow-route-leaking;     timeout <i>seconds</i>; } </pre>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],</p> <p>[edit protocols <i>isis</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</p>   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>  |
| <b>Description</b>         | <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"> <li>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.</li> <li>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.</li> <li>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.</li> </ol> |



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 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's 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.



## passive (Protocols IS-IS)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre> passive {     remote-node-id <i>address</i>;     remote-node-iso <i>iso-id</i>; } </pre>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <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 <b>interface</b> <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 <b>interface</b> <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 <b>interface</b> <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>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p><b>remote-node-id</b> <i>address</i> option introduced in Junos OS Release 14.2.</p> <p><b>remote-node-iso</b> <i>iso-id</i> option introduced in Junos OS Release 14.2.</p>   |
| <b>Description</b>         | <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 <b>interface</b> 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 <b>disable</b> statement at those hierarchy levels. The three states—enabling, disabling, or not running IS-IS on an interface—are mutually exclusive.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> Configuring IS-IS on a loopback interface automatically renders it as a passive interface, irrespective of whether the <b>passive</b> statement was used in the configuration of the interface.</p> </div> <p>If neither passive mode nor the <b>family iso</b> 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 <b>family iso</b> option at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.)</p> |
| <b>Default</b>             | By default, IS-IS must be configured on an interface or a level for direct interface addresses to be advertised into that level.  |



**Options**    **remote-node-id *address***—IP address of the remote link.

**remote-node-iso *iso-id***—ISO ID of the remote node.




**NOTE:** The options **remote-node-id *address*** and **remote-node-iso *iso-id*** do not apply under the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level.

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

**Related Documentation**    • [Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding on page 17](#)  
   • *disable*



## point-to-point

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | point-to-point;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| <b>Description</b>              | <p>Configure an IS-IS interface to behave like a point-to-point connection.</p> <p>You can use the <b>point-to-point</b> 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 <b>point-to-point</b> 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> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 20px;"> <p> <b>BEST PRACTICE:</b> On QFX10000 switches, we strongly recommend that you configure all IS-IS interfaces, including peer interfaces, as point-to-point interfaces. If you do not, you might experience session flaps, that is, IS-IS sessions that go down and then come back up, when IS-IS is configured in virtual routing and forwarding (VRF) instances. When you scale IS-IS in any scenario, you might also experience scaling issues if you do not configure IS-IS interfaces as point-to-point interfaces.</p> </div> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">IS-IS Overview on page 3</a></li> <li>• <a href="#">Example: Configuring Synchronization Between IS-IS and LDP</a></li> <li>• <a href="#">Understanding LDP-IGP Synchronization</a></li> <li>• <a href="#">Understanding IS-IS Designated Routers on page 71</a></li> </ul>  |



## preference (Protocols IS-IS)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>preference <i>preference</i>;</code>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>],</p> <p>[edit protocols isis <a href="#">level level-number</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>]</p>   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>  |
| <b>Description</b>              | <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 <b>preference</b> statement (for internal routes) or the <b>external-preference</b> statement.</p>  |
| <b>Options</b>                  | <p><b><i>preference</i></b>—Preference value.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 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>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Route Preferences Overview</a></li> <li>• <a href="#">Example: Configuring a Routing Policy to Redistribute BGP Routes with a Specific Community Tag into IS-IS</a></li> <li>• <a href="#">Example: Redistributing OSPF Routes into IS-IS on page 28</a></li> <li>• <a href="#">Understanding Routing Policies</a></li> <li>• <a href="#">Understanding BGP Communities, Extended Communities, and Large Communities as Routing Policy Match Conditions</a></li> <li>• <a href="#">external-preference on page 136</a></li> </ul> |



## prefix-export-limit (Protocols IS-IS)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>prefix-export-limit <i>number</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>],<br/> [edit protocols isis <a href="#">level level-number</a>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>]</p>   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>  |
| <b>Description</b>              | <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 <b>prefix-export-limit</b> statement. The <b>prefix-export-limit</b> 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 <a href="#">clear isis overload</a> command.</p> <p>The <a href="#">show isis overview</a> command displays the prefix export limit when it is configured.</p> |
| <b>Options</b>                  | <p><b><i>number</i></b>—Prefix limit.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> None</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Redistributing OSPF Routes into IS-IS on page 28</a></li> <li>• <a href="#">Example: Configuring a Routing Policy to Redistribute BGP Routes with a Specific Community Tag into IS-IS</a></li> <li>• <a href="#">Understanding BGP Communities, Extended Communities, and Large Communities as Routing Policy Match Conditions</a></li> </ul>  |



- *Understanding Routing Policies*

## prefix-segment

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>prefix-segment {   index <i>index</i>;   node-segment; }</pre>   |
| <b>Hierarchy Level</b>          | [edit policy-options policy-statement <i>policy-name</i> term <i>term-name</i> then]  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 17.2 for the MX Series routers, PTX Series routers, QFX5100 switches, and QFX10000 switches.</p> <p>Statement introduced in Junos OS Release 17.3R1 for the QFX5110 and QFX5200 switches.</p>   |
| <b>Description</b>              | <p>Configure prefix segment attributes such as index and node segment. Prefix segment index is the index assigned to a specific prefix. This is used by all other remote routers in the network to index the prefix into respective segment routing global blocks (SRGB) to derive the segment identifier and to forward the traffic destined for the current prefix. Prefix segment index is provisioned through policy for each prefix along with the option to mark it as a node segment.</p> <p><b>index <i>index</i></b>—Specify the prefix segment index<br/> <b>Range:</b> 0 through 199,999</p> <p><b>node-segment</b>—Set node segment flag for this prefix segment.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Support for SRGB, Anycast, and Prefix Segments in SPRING for IS-IS Protocol on page 114</a></li> <li>• <a href="#">Configuring Anycast and Prefix segments in SPRING for IS-IS Protocol on page 116</a></li> <li>• <a href="#">adjacency-segment on page 125</a></li> </ul>  |



## priority (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>  |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code><br><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><br><code>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</code>   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| <b>Description</b>              | <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> |
| <b>Options</b>                  | <b><i>number</i></b> —Priority value.<br><b>Range:</b> 0 through 127<br><b>Default:</b> 64  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS Designated Routers on page 71</a></li></ul>  |



## reference-bandwidth (Protocols IS-IS)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>reference-bandwidth <i>reference-bandwidth</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],<br/>         [edit protocols <i>isis</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/>         Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>         Statement introduced in Junos OS Release 12.1 for the QFX Series.<br/>         Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>   |
| <b>Description</b>              | <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. Tweak the reference bandwidth to route traffic across the fastest interface.</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. By default, IS-IS has a metric of 10. Modify the configured reference bandwidth to increase the cost of an interface in order to avoid a slow route or you can modify the reference bandwidth value to decrease the cost of a preferred interface. A route that has a lower cost is selected over a high cost route. Note that fractional metric values are not allowed. Therefore, if you configure a reference bandwidth that is higher than the actual bandwidth, which results in a fractional metric value then IS-IS takes the metric value as 1 because fractional metric values are not allowed.</p> |
| <b>Options</b>                  | <p><i>reference-bandwidth</i>—Reference bandwidth value in bits per second.<br/> <b>Range:</b> 9600 through 1,000,000,000,000 bps<br/> <b>Default:</b> None</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/>         routing-control—To add this statement to the configuration.</p>  |



- Related Documentation**
- *Understanding IS-IS Configuration*
  - [Example: Configuring IS-IS on page 11](#)
  - [https://www.juniper.net/us/en/training/certification/JNCIP\\_studyguide.pdf](https://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf)

---

## rib-group (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>rib-group {<br/>    inet <i>group-name</i>;<br/>    inet6 <i>group-name</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| <b>Description</b>              | <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> |
| <b>Options</b>                  | <p><b><i>group-name</i></b>—Name of the routing table group.</p> <p><b>inet</b>—Install IPv4 IS-IS routes.</p> <p><b>inet6</b>—Install IPv6 IS-IS routes.</p>   |
| <b>Required Privilege Level</b> | <p><b>routing</b>—To view this statement in the configuration.</p> <p><b>routing-control</b>—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li><li>• <i>Example: Importing Direct and Static Routes Into a Routing Instance</i></li><li>• <i>Configuring BGP Routing Table Groups</i></li></ul>  |



## source-packet-routing (Protocols IS-IS and OSPF)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre> source-packet-routing {     adjacency-segment {         hold-time <i>hold-time</i>;     }     disable;     explicit-null;     node-segment {         index-range <i>index range</i>;         ipv4-index <i>index</i>;         ipv6-index <i>index</i>;     }     srgb {         start-label <i>start-label</i>;         index-range <i>index range</i>;     } } </pre>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],<br/> [edit protocols <i>isis</i>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]<br/> [edit logical-systems <i>logical-system-name</i> protocols <i>isis</i> <i>level level-number</i>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i> <i>level level-number</i>],<br/> [edit protocols <i>isis</i> <i>level level-number</i>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i> <i>level level-number</i>]<br/> [edit logical-systems <i>logical-system-name</i> protocols ospf],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf],<br/> [edit protocols ospf],<br/> [edit routing-instances <i>routing-instance-name</i> protocols ospf]<br/> [edit logical-systems <i>logical-system-name</i> protocols ospf <i>area-number</i>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <i>area-number</i>],<br/> [edit protocols ospf <i>area-number</i>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols ospf <i>area-number</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 16.2R1 for MX Series routers.<br/> Statement introduced in Junos OS Release 16.1X65 for the PTX1000 routers.<br/> Statement introduced in Junos OS Release 17.2R1 for QFX5100 switches and QFX10000 switches.<br/> Support for <b>explicit-null</b> and <b>adjacency-segment</b> statements introduced in Junos OS Release 17.2R1 for MX Series routers, PTX Series routers, QFX5100 switches and QFX10000 switches.<br/> Statement introduced in Junos OS Release 17.3R1 for QFX5110 and QFX5200 switches.</p>   |
| <b>Description</b>         | <p>Enable source packet routing in networking (SPRING) feature on IS-IS levels or OSPF areas.</p>   |





**NOTE:** Source packet routing for IPv6 is supported only for IS-IS. The `adjacency-segment`, `explicit-null`, and `srgb` statements are supported only for IS-IS.

The remaining statements are explained separately.

**Default** Disabled.

**Options** **disable**—Disable source packet routing from a specific level.  
**explicit-null**—Configure E and P bits in all prefix segment identifier (SID) advertisements.

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

**Related Documentation**

- [Understanding Source Packet Routing in Networking \(SPRING\) on page 111](#)
- [Support for SRGB, Anycast, and Prefix Segments in SPRING for IS-IS Protocol on page 114](#)
- [Configuring Anycast and Prefix segments in SPRING for IS-IS Protocol on page 116](#)



## source-packet-routing-disable (Protocols IS-IS)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | source-packet-routing {<br>disable;<br>}   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ]         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 16.2R1 for MX Series routers.<br>Statement introduced in Junos OS Release 16.1X65 for the PTX1000 routers.<br>Statement introduced in Junos OS Release 17.2R1 for QFX5100, QFX5110, and QFX10000 switches.<br>Statement introduced in Junos OS Release 17.3R1 for QFX5110 and QFX5200 switches.   |
| <b>Description</b>              | Disable source packet routing in networking (SPRING) feature at a specific level.  |
| <b>Default</b>                  | Enabled, if <a href="#">source-packet-routing</a> is configured globally at:<br><br>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i><br>protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ] |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">isis on page 157</a></li> </ul>   |



## spf-options (Protocols IS-IS)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre>spf-options {     delay <i>milliseconds</i>;     holddown <i>milliseconds</i>;     rapid-runs <i>number</i>; }</pre>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>],</p> <p>[edit protocols <a href="#">isis</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>]</p>   |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>   |
| <b>Description</b>         | <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"> <li>• The delay in the time between the detection of a topology change and when the SPF algorithm actually runs.</li> <li>• The maximum number of times that the SPF algorithm can run in succession before the hold-down timer begins.</li> <li>• 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.</li> </ul> <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 <b>delay</b> and <b>rapid-runs</b> statements.</p> |
| <b>Options</b>             | <p><b>delay <i>milliseconds</i></b>—Time interval between the detection of a topology change and when the SPF algorithm runs.</p> <p><b>Range:</b> 50 through 1000 milliseconds</p> <p><b>Default:</b> 200 milliseconds</p> <p><b>holddown <i>milliseconds</i></b>—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><b>Range:</b> 2000 through 10,000 milliseconds</p>  |



**Default:** 5000 milliseconds

**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

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.        |
|                                 | routing-control—To add this statement to the configuration. |

|                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | • <a href="#">Understanding Loop-Free Alternate Routes for IS-IS on page 95</a>                          |
|                              | • <a href="#">Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 99</a> |



## srgb

---

|                          |  |
|--------------------------|--|
| Syntax                   | <pre>srgb {<br/>  start-label <i>start-label-value</i>;<br/>  index-range <i>index range-value</i>;<br/>}</pre>  |
| Hierarchy Level          | [edit logical-systems <i>name</i> protocols isis <a href="#">source-packet-routing</a> ],<br>[edit logical-systems <i>name</i> routing-instances <i>name</i> protocols isis <a href="#">source-packet-routing</a> ],<br>[edit protocols isis <a href="#">source-packet-routing</a> ],<br>[edit routing-instances <i>name</i> protocols isis <a href="#">source-packet-routing</a> ]  |
| Release Information      | Statement introduced in Junos OS Release 17.2 for the MX Series routers in Enhanced-IP mode, PTX Series routers, QFX5100 switches, and QFX10000 switches.<br>Statement introduced in Junos OS Release 17.3 for QFX5110 and QFX5200 switches.   |
| Description              | Configure the segment routing global block (SRGB) in source packet routing in networking (SPRING) or segment routing (SR). The SRGB label range is based on the start label and the index range. The value of the start label indicates the start of the label range, and the value of the index range along with the value of the start label indicate the end of the label range.  |
| Options                  | <b>index-range <i>index-range-value</i></b> —Index range of the SRGB label block.<br><b>Range:</b> 32 through 1,048,559<br><br><b>start-label <i>start-label-value</i></b> —Start of the SRGB label block.<br><b>Range:</b> 16 through 1,048,575   |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Support for SRGB, Anycast, and Prefix Segments in SPRING for IS-IS Protocol on page 114</a></li><li>• <i>Configuring Segment Routing Global Blocks Label Ranges in SPRING for IS-IS Protocol</i></li><li>• <i>Example: Configuring Segment Routing Global Blocks in SPRING for IS-IS to Increase Network Speed</i></li><li>• <a href="#">source-packet-routing on page 195</a></li></ul> |



## topologies (Protocols IS-IS)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> topologies {   ipv4-multicast;   ipv6-multicast;   ipv6-unicast; } </pre>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>],</p> <p>[edit protocols <a href="#">isis</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>]</p>                          |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>   |
| <b>Description</b>              | <p>Configure alternate IS-IS topologies.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li> <li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 47</a></li> <li>• <a href="#">IS-IS Multicast Topologies Overview on page 46</a></li> <li>• <a href="#">Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 37</a></li> </ul> |



## traceoptions (Protocols IS-IS)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>name</i> &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>isis</b>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>isis</b>],<br/>         [edit protocols <b>isis</b>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols <b>isis</b>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.<br/>         Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>         Statement introduced in Junos OS Release 12.1 for the QFX Series.<br/>         Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>  |
| <b>Description</b>         | <p>Configure IS-IS protocol-level tracing options. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>  |



**NOTE:** The **traceoptions** statement is not supported on QFabric systems.

|                |  |
|----------------|--|
| <b>Default</b> | <p>The default IS-IS protocol-level tracing options are those inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.</p>   |
| <b>Options</b> | <p><b>disable</b>—(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 <b>all</b>.</p> <p><b>file <i>name</i></b>—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 <b>/var/log</b>. We recommend that you place IS-IS tracing output in the file <b>isis-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 10 files</p> |



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

#### 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
- **layer2-map**—Mapped ARP or neighbor discovery next hops in the kernel
- **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.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.<br>routing-control and trace-control—To add this statement to the configuration. |
|---------------------------------|---|

|                              |   |
|------------------------------|---|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces</i></li><li>• <a href="#">Example: Enabling Packet Checksums on IS-IS Interfaces for Error Checking on page 72</a></li><li>• <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i></li><li>• <i>Understanding Link-State PDU Throttling for IS-IS Interfaces</i></li><li>• <i>Understanding Checksums on IS-IS Interfaces for Error Checking</i></li><li>• <i>Understanding the Transmission Frequency for CSNPs on IS-IS Interfaces</i></li></ul> |
|------------------------------|---|



## traffic-engineering (Protocols IS-IS)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre> traffic-engineering {   disable;   credibility-protocol-preference;   igp-topology;   family inet {     shortcuts {       multicast-rpf-routes;     }   }   family inet-mpls {     shortcuts;   }   family inet6 {     shortcuts;   }   family inet6-mpls {     shortcuts;   }   multipath {     lsp-equal-cost;   } } </pre>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>isis</b>],<br/> [edit protocols <b>isis</b>]</p>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.<br/> Support for the <b>family</b> statement introduced in Junos OS Release 9.3.<br/> Support for the <b>credibility-protocol-preference</b> statement introduced in Junos OS Release 9.4.<br/> Support for the <b>multipath</b> statement introduced in Junos OS Release 9.6.<br/> Support for the <b>lsp-equal-cost</b> statement introduced in Junos OS Release 9.6.<br/> Statement introduced in Junos OS Release 12.1 for the QFX Series.<br/> Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.<br/> Support for <b>inet-mpls</b> and <b>inet6-mpls</b> statements introduced in Junos OS Release 17.2 for MX Series, PTX Series, QFX5100 switches, and QFX10000 line switches.<br/> Support for <b>inet-mpls</b> and <b>inet6-mpls</b> statements introduced in Junos OS Release 17.3 for QFX5110 and QFX5200 switches.<br/> Support for <b>igp-topology</b> statement introduced in Junos OS Release 17.4R1 for MX Series, and PTX Series.</p> |
| <b>Description</b>         | <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</p>  |



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.

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.

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 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 IPv4 MPLS or IPv6 MPLS shortcuts explicitly for segment routing, include the **inet-mpls** statement for IPv4 MPLS traffic and the **inet6-mpls** statement for IPv6 MPLS traffic.

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.

**igp-topology**—Download IGP topology information into the traffic engineering database (TED). In Junos OS, the IGPs install topology information into a database called the traffic engineering database. The traffic engineering database contains the aggregated topology information. The IGP routes are installed by the traffic engineering database on behalf of the corresponding IGP into a user-visible routing table called `Isdist.0`, subject to route policies.

The remaining statements are explained separately. See [CLI Explorer](#).

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | 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*
  - *traffic-engineering (OSPF)*
  - *Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts*

## use-source-packet-routing (Protocols IS-IS)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | use-source-packet-routing;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis backup-spf-options],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis backup-spf-options],<br>[edit protocols isis backup-spf-options],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis backup-spf-options] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 16.2 for the M, MX, and PTX Series.<br>Statement introduced in Junos OS Release 17.2R1 for QFX5100 and QFX10000 switches.<br>Statement introduced in Junos OS Release 17.3R1 for QFX5110 and QFX5200 switches.   |
| <b>Description</b>              | Enable use of source packet routing node segment labels for computing backup paths for normal IPv4 or IPv6 IS-IS prefixes and primary IS-IS source packet routing node segments.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">isis on page 157</a></li> <li>• <a href="#">source-packet-routing on page 195</a></li> </ul>   |



## wide-metrics-only

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | wide-metrics-only;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <b>level</b> <i>level-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>level</b> <i>level-number</i> ],<br>[edit protocols isis <b>level</b> <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>level</b> <i>level-number</i> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>              | <p>Configure IS-IS to generate metric values greater than 63 on a per IS-IS level basis.</p> <p>Normally, IS-IS metrics can have values up to 63, and IS-IS generates two type, length, and value (TLV) tuples, one for an IS-IS adjacency and the second for an IP prefix. To allow IS-IS to support traffic engineering, a second pair of TLVs has been added to IS-IS, one for IP prefixes and the second for IS-IS adjacency and traffic engineering information. With these TLVs, IS-IS metrics can have values up to 16,777,215 (<math>2^{24} - 1</math>).</p> <p>To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, include the <b>wide-metrics-only</b> statement.</p> |
| <b>Default</b>                  | 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.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>te-metric</i></li><li>• <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i></li><li>• <i>Understanding Wide IS-IS Metrics for Traffic Engineering</i></li></ul>   |



## CHAPTER 8

# Operational Commands

- clear isis adjacency
- clear isis database
- clear isis overload
- clear isis statistics
- show isis adjacency
- show isis authentication
- show isis backup coverage
- show isis backup label-switched-path
- show isis backup spf results
- show isis database
- show isis hostname
- show isis interface
- show isis overview
- show isis route
- show isis spf
- show isis statistics



## clear isis adjacency

---

|   |   |
|---|---|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 212</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 212</a>  |
| <b>Syntax</b>                                     | <pre>clear isis adjacency &lt;all&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;neighbor&gt;</pre>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>clear isis adjacency &lt;all&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;neighbor&gt;</pre>  |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br><b>all</b> option introduced in Junos OS Release 14.2.  |
| <b>Description</b>                                | Remove entries from the IS-IS adjacency database.   |
| <b>Options</b>                                    | <p><b>all</b>—Remove all entries from the adjacency database.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear all adjacencies for the specified routing instance only.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear all adjacencies for the specified interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear adjacencies for the specified neighbor only.</p> |
| <b>Required Privilege Level</b>                   | clear   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">show isis adjacency on page 220</a></li></ul>   |
| <b>List of Sample Output</b>                      | <a href="#">clear isis adjacency on page 213</a><br><a href="#">clear isis adjacency all on page 213</a>  |
| <b>Output Fields</b>                              | See <a href="#">show isis adjacency</a> 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 karaku1

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State      Hold (secs) SNPA
so-1/0/0.0     karaku1         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 adjacency all

```
user@host> clear isis adjacency all
IS-IS adjacency database:
Interface      System          L State      Hold (secs) SNPA
so-1/0/0.0     karaku1         3 Initializing 26
so-1/1/3.0     1921.6800.5080 3 Initializing 24
so-5/0/0.0     1921.6800.5080 3 Initializing 21
```



## clear isis database

---

|   |   |
|---|---|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 214</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 214</a>  |
| <b>Syntax</b>                                     | <pre>clear isis database &lt;all&gt; &lt;entries&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>clear isis database &lt;all&gt; &lt;entries&gt; &lt;instance <i>instance-name</i>&gt;</pre>  |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.<br>Command introduced in 15.1X53-D30 for QFX10002 switch.  |
| <b>Description</b>                                | Remove the entries from the IS-IS link-state database, which contains prefixes and topology information.  |
| <b>Options</b>                                    | <b>all</b> —Remove all entries from the IS-IS link-state database for all routing instances.<br><br><b>entries</b> —(Optional) Name of the database entry.<br><br><b>instance <i>instance-name</i></b> —(Optional) Clear all entries for the specified routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | clear   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">show isis database on page 235</a></li></ul>  |
| <b>List of Sample Output</b>                      | <a href="#">clear isis database on page 215</a>   |
| <b>Output Fields</b>                              | See <a href="#">show isis database</a> 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 all** 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 all

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

---

|   |   |
|---|---|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 216</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 216</a>  |
| <b>Syntax</b>                                     | <pre>clear isis overload &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>clear isis overload &lt;instance <i>instance-name</i>&gt;</pre>  |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| <b>Description</b>                                | <p>Reset the IS-IS dynamic overload bit. This command can appear to not work, continuing to display <b>overload</b> after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.</p> <p>When other routers detect that the overload bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes.</p> |
| <b>Options</b>                                    | <p><b>none</b>—Reset the IS-IS dynamic overload bit.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>  |
| <b>Required Privilege Level</b>                   | clear   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">show isis database on page 235</a></li></ul>  |
| <b>List of Sample Output</b>                      | <a href="#">clear isis overload on page 217</a>   |
| <b>Output Fields</b>                              | See <a href="#">show isis database</a> 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

---

|  |   |
|--|---|
| List of Syntax                             | <a href="#">Syntax on page 218</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 218</a>  |
| Syntax                                     | <code>clear isis statistics</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>  |
| Syntax (EX Series Switches and QFX Series) | <code>clear isis statistics</code><br><code>&lt;instance <i>instance-name</i>&gt;</code>  |
| Release Information                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  |
| Description                                | Set statistics about IS-IS traffic to zero.   |
| Options                                    | <b>none</b> —Set IS-IS traffic statistics to zero for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(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"><li>• <a href="#">show isis statistics on page 270</a></li></ul>  |
| List of Sample Output                      | <a href="#">clear isis statistics on page 218</a>   |
| Output Fields                              | See <a href="#">show isis statistics</a> 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



## show isis adjacency

---

|   |  |
|---|--|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 220</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 220</a>   |
| <b>Syntax</b>                                     | <pre>show isis adjacency &lt;system-id&gt; &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis adjacency &lt;system-id&gt; &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt;</pre>  |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>                                | Display information about IS-IS neighbors.   |
| <b>Options</b>                                    | <p><b>none</b>—Display standard information about IS-IS neighbors for all routing instances.</p> <p><b><i>system id</i></b>—(Optional) Display information about IS-IS neighbors for the specified intermediate system.</p> <p><b>brief   detail   extensive</b>—(Optional) Display standard information about IS-IS neighbors with the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about IS-IS neighbors for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Display information about IS-IS neighbors for all logical systems or for a particular logical system.</p> |
| <b>Required Privilege Level</b>                   | view   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">clear isis adjacency on page 212</a></li></ul>   |
| <b>List of Sample Output</b>                      | <a href="#">show isis adjacency on page 223</a><br><a href="#">show isis adjacency brief on page 223</a><br><a href="#">show isis adjacency detail on page 224</a><br><a href="#">show isis adjacency extensive on page 224</a>  |



**Output Fields** Table 6 on page 221 describes the output fields for the **show isis adjacency** command. Output fields are listed in the approximate order in which they appear.

*Table 6: show isis adjacency Output Fields*

| Field Name                                | Field Description   | Level of Output         |
|---|---|-------------------------|
| <b>Interface</b>                          | Interface through which the neighbor is reachable.  | All levels              |
| <b>System</b>                             | System identifier ( <b>sysid</b> ), displayed as a name, if possible.   | <b>brief</b>            |
| <b>L or Level</b>                         | Level: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul> An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address. | All levels              |
| <b>State</b>                              | State of the adjacency: <b>Up</b> , <b>Down</b> , <b>New</b> , <b>One-way</b> , <b>Initializing</b> , or <b>Rejected</b> .  | All levels              |
| <b>Hold (secs)</b>                        | Remaining hold time of the adjacency.   | <b>brief</b>            |
| <b>SNPA</b>                               | Subnetwork point of attachment (MAC address of the next hop).   | <b>brief</b>            |
| <b>Expires in</b>                         | How long until the adjacency expires, in seconds.   | <b>detail</b>           |
| <b>Priority</b>                           | Priority to become the designated intermediate system.  | <b>detail extensive</b> |
| <b>Up/Down transitions</b>                | Count of adjacency status changes from <b>Up</b> to <b>Down</b> or from <b>Down</b> to <b>Up</b> .  | <b>detail</b>           |
| <b>Last transition</b>                    | Time of the last <b>Up/Down</b> transition.   | <b>detail</b>           |
| <b>Circuit type</b>                       | Bit mask of levels on this interface: 1=Level 1 router; 2=Level 2 router; 3=both Level 1 and Level 2 router.  | <b>detail</b>           |
| <b>Speaks</b>                             | Protocols supported by this neighbor.   | <b>detail extensive</b> |
| <b>MAC address</b>                        | MAC address of the interface.   | <b>detail extensive</b> |
| <b>Topologies</b>                         | Supported topologies.   | <b>detail extensive</b> |
| <b>Restart capable</b>                    | Whether a neighbor is capable of graceful restart: <b>Yes</b> or <b>No</b> .  | <b>detail extensive</b> |
| <b>Adjacency advertisement: Advertise</b> | This routing device has signaled to advertise this interface to its neighbors in their link-state PDUs.   | <b>detail extensive</b> |
| <b>Adjacency advertisement: Suppress</b>  | This neighbor has signaled not to advertise the interface in the routing device's outbound link-state PDUs.   | <b>detail extensive</b> |



*Table 6: show isis adjacency Output Fields (continued)*

| Field Name           | Field Description                               | Level of Output         |
|----------------------|---|-------------------------|
| <b>IP addresses</b>  | IP address of this neighbor.                    | <b>detail extensive</b> |
| IPv6 Address         | IPv6 address of the neighbor.                   | <b>detail extensive</b> |
| Level 1 IPv4 Adj-SID | Level 1 IPv4 node-SID of the adjacent neighbor. | <b>detail extensive</b> |
| Level 1 IPv6 Adj-SID | Level 1 IPv6 node-SID of the adjacent neighbor. |                         |
| Level 2 IPv4 Adj-SID | Level 2 IPv4 node-SID of the adjacent neighbor. | <b>detail extensive</b> |
| Level 2 IPv6 Adj-SID | Level 2 IPv6 node-SID of the adjacent neighbor. | <b>detail extensive</b> |



Table 6: *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"> <li>• <b>When</b>—Time at which an IS-IS adjacency transition occurred.</li> <li>• <b>State</b>—Current state of the IS-IS adjacency (<b>up</b>, <b>down</b>, or <b>rejected</b>). <ul style="list-style-type: none"> <li>• <b>Up</b>—Adjacency is up and operational.</li> <li>• <b>Down</b>—Adjacency is down and not available.</li> <li>• <b>Rejected</b>—Adjacency has been rejected.</li> </ul> </li> <li>• <b>Event</b>—Type of transition that occurred. <ul style="list-style-type: none"> <li>• <b>Seenself</b>—Possible routing loop has been detected.</li> <li>• <b>Interface down</b>—IS-IS interface has gone down and is no longer available.</li> <li>• <b>Error</b>—Adjacency error.</li> </ul> </li> <li>• <b>Down reason</b>—Reason that an IS-IS adjacency is down: <ul style="list-style-type: none"> <li>• <b>3-Way Handshake Failed</b>—Connection establishment failed.</li> <li>• <b>Address Mismatch</b>—Address mismatch caused link failure.</li> <li>• <b>Aged Out</b>—Link expired.</li> <li>• <b>ISO Area Mismatch</b>—IS-IS area mismatch caused link failure.</li> <li>• <b>Bad Hello</b>—Unacceptable hello message caused link failure.</li> <li>• <b>BFD Session Down</b>—Bidirectional failure detection caused link failure.</li> <li>• <b>Interface Disabled</b>—IS-IS interface is disabled.</li> <li>• <b>Interface Down</b>—IS-IS interface is unavailable.</li> <li>• <b>Interface Level Disabled</b>—IS-IS level is disabled.</li> <li>• <b>Level Changed</b>—IS-IS level has changed on the adjacency.</li> <li>• <b>Level Mismatch</b>—Levels on adjacency are not compatible.</li> <li>• <b>MPLS LSP Down</b>—Label-switched path (LSP) is unavailable.</li> <li>• <b>MT Topology Changed</b>—IS-IS topology has changed.</li> <li>• <b>MT Topology Mismatch</b>—IS-IS topology is mismatched.</li> <li>• <b>Remote System ID Changed</b>—Adjacency peer system ID changed.</li> <li>• <b>Protocol Shutdown</b>—IS-IS protocol is disabled.</li> <li>• <b>CLI Command</b>—Adjacency brought down by user.</li> <li>• <b>Unknown</b>—Unknown.</li> </ul> </li> </ul> | 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 223](#).



### 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, IPV6-Unicast Restart capable: Yes, Adjacency advertisement:
  Advertise
  LAN id: pro-bng3-c-F.02, IP addresses: 11.1.1.2
  IPv6 addresses: fe80::2a0:a514:0:4745
  Level 1 IPv4 Adj-SID: 299808, IPv6 Adj-SID: 299824
```

### 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, IPV6-Unicast
  Restart capable: Yes, Adjacency advertisement: Advertise
  IP addresses: 11.1.1.2
  IPv6 addresses: fe80::2a0:a514:0:3e45
  Level 1 IPv4 Adj-SID: 300112, IPv6 Adj-SID: 300304
  Level 2 IPv4 Adj-SID: 300320, IPv6 Adj-SID: 300336
  Transition log:
  When                State      Event      Down reason
  Thu Mar 26 06:13:18  Up        Seenself
```



## show isis authentication

|   |   |
|---|---|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 225</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 225</a>  |
| <b>Syntax</b>                                     | <pre>show isis authentication &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis authentication &lt;instance <i>instance-name</i>&gt;</pre>   |
| <b>Release Information</b>                        | <p>Command introduced in Junos OS Release 7.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for hitless authentication key rollover introduced in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| <b>Description</b>                                | Display information about Intermediate System-to-Intermediate System (IS-IS) authentication.  |
| <b>Options</b>                                    | <p><b>none</b>—Display information about IS-IS authentication.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display IS-IS authentication for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>                        |
| <b>Required Privilege Level</b>                   | view  |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li><a href="#">show security keychain</a></li> </ul>  |
| <b>List of Sample Output</b>                      | <a href="#">show isis authentication on page 226</a><br><a href="#">show isis authentication (With Hitless Authentication Key Rollover Configured) on page 226</a>  |
| <b>Output Fields</b>                              | <p><a href="#">Table 7 on page 225</a> describes the output fields for the <b>show isis authentication</b> command. Output fields are listed in the approximate order in which they appear.</p>   |

*Table 7: show isis authentication Output Fields*

| Field Name       | Field Description |
|------------------|-------------------|
| <b>Interface</b> | Interface name.   |



*Table 7: show isis authentication Output Fields (continued)*

| Field Name                   | Field Description   |
|------------------------------|---|
| <b>Level</b>                 | IS-IS level.  |
| <b>IIH Auth</b>              | IS-IS Hello (IIH) packet authentication type.<br><br>Displays the name of the active keychain if hitless authentication key rollover is configured. |
| <b>CSN Auth</b>              | Complete sequence number authentication type.   |
| <b>PSN Auth</b>              | Partial sequence number authentication type.  |
| <b>L1 LSP Authentication</b> | Layer 1 link-state PDU authentication type.   |
| <b>L2 LSP Authentication</b> | 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

|   |   |
|---|---|
| <b>Syntax</b>                                     | <pre>show isis backup coverage &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis backup coverage &lt;instance <i>instance-name</i>&gt;</pre>  |
| <b>Release Information</b>                        | <p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>   |
| <b>Description</b>                                | Display information about the level of backup coverage available.   |
| <b>Options</b>                                    | <p><b>none</b>—Display information about the level of backup coverage available for all the nodes and prefixes in the network.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the level of backup coverage for a specific IS-IS routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                   | view  |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Loop-Free Alternate Routes for IS-IS on page 95</a></li> <li>• <a href="#">Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 99</a></li> <li>• <a href="#">show isis backup label-switched-path on page 229</a></li> </ul>   |
| <b>List of Sample Output</b>                      | <a href="#">show isis backup coverage on page 228</a>   |
| <b>Output Fields</b>                              | <p><a href="#">Table 8 on page 227</a> lists the output fields for the <b>show isis backup coverage</b> command. Output fields are listed in the approximate order in which they appear.</p>  |

*Table 8: show isis backup coverage Output Fields*

| Field Name | Field Description  |
|------------|--|
| Topology   | Type of topology or address family: <b>IPv4 Unicast</b> or <b>IPv6 Unicast</b> . |



*Table 8: show isis backup coverage Output Fields (continued)*

| Field Name   | Field Description  |
|--------------|--|
| <b>Level</b> | IS-IS level: <ul style="list-style-type: none"> <li>• 1—Level 1</li> <li>• 2—Level 2</li> </ul>              |
| <b>Node</b>  | By topology, the percentage of all routes configured on the node that are protected through backup coverage. |
| <b>IPv4</b>  | Percentage of IPv4 unicast routes that are protected through backup coverage.                                |
| <b>IPv6</b>  | Percentage of IPv6 unicast routes that are protected through backup coverage.                                |
| <b>CLNS</b>  | 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    Level1  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

|   |  |
|---|--|
| <b>Syntax</b>                                     | <b>show isis backup label-switched-path</b><br><b>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</b>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <b>show isis backup label-switched-path</b>  |
| <b>Release Information</b>                        | Command introduced in Junos OS Release 9.5.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| <b>Description</b>                                | Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes.   |
| <b>Options</b>                                    | <b>none</b> —Display information about MPLS LSPs designated as backup routes for IS-IS routes.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  |
| <b>Required Privilege Level</b>                   | view   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Loop-Free Alternate Routes for IS-IS on page 95</a></li> <li>• <a href="#">Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 99</a></li> <li>• <a href="#">show isis backup coverage on page 227</a></li> </ul> |
| <b>List of Sample Output</b>                      | <a href="#">show isis backup label-switched-path on page 230</a>   |
| <b>Output Fields</b>                              | <a href="#">Table 9 on page 229</a> lists the output fields for the <b>show isis backup label-switched-path</b> command. Output fields are listed in the approximate order in which they appear.   |

*Table 9: show isis backup label-switched-path Output Fields*

| Field Name              | Field Description  |
|-------------------------|--|
| <b>Backup MPLS LSPs</b> | List of MPLS LSPs designated as backup paths for IS-IS routes. |
| <b>Egress</b>           | IP address of the egress routing device for the LSP.           |



Table 9: show isis backup label-switched-path Output Fields (continued)

| Field Name  | Field Description  |
|-------------|--|
| Status      | <p>State of the LSP:</p> <ul style="list-style-type: none"> <li>• <b>Up</b>—The routing device can detect RSVP hello messages from the neighbor.</li> <li>• <b>Down</b>—The routing device has received one of the following indications: <ul style="list-style-type: none"> <li>• Communication failure from the neighbor.</li> <li>• Communication from IGP that the neighbor is unavailable.</li> <li>• Change in the sequence numbers in the RSVP hello messages sent by the neighbor.</li> </ul> </li> <li>• <b>Deleted</b>—LSP is no longer available as a backup path.</li> </ul> |
| 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

|                                    |  |
|------------------------------------|--|
| <b>Syntax</b>                      | <pre>show isis backup spf results &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;no-coverage&gt; &lt;topology (ipv4-unicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>  |
| <b>Syntax (EX Series Switches)</b> | <pre>show isis backup spf results &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;no-coverage&gt; &lt;topology (ipv4-unicast   unicast)&gt;</pre>  |
| <b>Release Information</b>         | Command introduced in Junos OS Release 9.5.  |
| <b>Description</b>                 | Display information about IS-IS shortest-path-first (SPF) calculations for backup paths.   |
| <b>Options</b>                     | <p><b>none</b>—Display information about IS-IS SPF calculations for all backup paths for all destination nodes.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display SPF calculations for backup paths for the specified routing instance.</p> <p><b>level (1   2)</b>—(Optional) Display SPF calculations for the backup paths for the specified IS-IS level.</p> <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Display SPF calculations for the backup paths for all logical systems or on a particular logical system.</p> <p><b>no-coverage</b>—(Optional) Display SPF calculations only for destinations that do not have backup coverage.</p> <p><b>topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)</b>—(Optional) Display SPF calculations for backup paths for the specified topology only.</p> |
| <b>Required Privilege Level</b>    | view   |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Link and Node Protection for IS-IS Routes on page 95</a></li> <li>• <a href="#">show isis backup coverage on page 227</a></li> <li>• <a href="#">Understanding Loop-Free Alternate Routes for IS-IS on page 95</a></li> <li>• <a href="#">Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 99</a></li> </ul>  |
| <b>List of Sample Output</b>       | <a href="#">show isis backup spf results on page 232</a>   |



[show isis backup spf results no-coverage on page 233](#)

**Output Fields** [Table 10 on page 232](#) lists the output fields for the **show isis backup spf results** command. Output fields are listed in the approximate order in which they appear.

*Table 10: show isis backup spf results Output Fields*

| Field Name              | Field Description  |
|-------------------------|--|
| <b>node-name</b>        | Name of the destination node.  |
| <b>Address</b>          | Address of the destination node.   |
| <b>Primary next-hop</b> | Interface and name of the node of the primary next hop to reach the destination.                   |
| <b>Root</b>             | Name of the next-hop neighbor.   |
| <b>Metric</b>           | Metric to the node.  |
| <b>Eligible</b>         | Indicates that the next-hop neighbor has been designated as a backup path to the destination node. |
| <b>Backup next-hop</b>  | Name of the interface of the backup next hop.  |
| <b>SNPA</b>             | Subnetwork point of attachment (MAC address of the next hop).                                      |
| <b>LSP</b>              | Name of the MPLS label-switched path (LSP) designated as a backup path.                            |
| <b>Not eligible</b>     | Indicates that the next-hop neighbor cannot function as a backup path to the destination.          |
| <b>Reason</b>           | Describes why the next-hop neighbor is designated as <b>Not eligible</b> as a backup path.         |

## Sample Output

### show isis backup spf results

```

user@host> show isis backup spf results

IS-IS level 1 SPF results:
  0 nodes

IS-IS level 2 SPF results:
banff.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
  Primary next-hop: so-6/0/0.0, IPV6, olympic
  Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
    Root: camaro, Root Metric: 10, Metric: 10
    Not eligible, Reason: Primary next-hop multipath
    Root: olympic, Root Metric: 10, Metric: 10

```



```

    Not eligible, Reason: Primary next-hop multipath
    Root: glacier, Root Metric: 10, Metric: 25
    Not eligible, Reason: Primary next-hop multipath
crater.00
    Primary next-hop: so-6/0/0.0, IPV4, olympic
    Primary next-hop: so-6/0/0.0, IPV6, olympic
    Root: olympic, Root Metric: 10, Metric: 10
    Not eligible, Reason: Primary next-hop link fate sharing
    Root: glacier, Root Metric: 10, Metric: 15
    Eligible, Backup next-hop: as0.0, IPV4, glacier
    Eligible, Backup next-hop: as0.0, IPV6, glacier
    Root: camaro, Root Metric: 10, Metric: 20
    Not eligible, Reason: Interface is already covered
olympic.00
    Primary next-hop: so-6/0/0.0, IPV4, olympic
    Primary next-hop: so-6/0/0.0, IPV6, olympic
    Root: olympic, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
    Root: camaro, Root Metric: 10, Metric: 20
    track-item: olympic.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
    Root: glacier, Root Metric: 10, Metric: 20
    track-item: olympic.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
camaro.00
    Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
    Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
    Root: camaro, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
    Root: glacier, Root Metric: 10, Metric: 20
    track-item: camaro.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
    Root: olympic, Root Metric: 10, Metric: 20
    track-item: camaro.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
glacier.00
    Primary next-hop: as0.0, IPV4, glacier
    Primary next-hop: as0.0, IPV6, glacier
    Root: glacier, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
    Root: camaro, Root Metric: 10, Metric: 20
    track-item: glacier.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
    Root: olympic, Root Metric: 10, Metric: 20
    track-item: glacier.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
5 nodes

```

#### show isis backup spf results no-coverage

```

user@host> show isis backup spf results no-coverage
IS-IS level 1 SPF results:
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:

```

olympic.00
Primary next-hop: so-6/0/0.0, IPV4, olympic
Primary next-hop: so-6/0/0.0, IPV6, olympic
  Root: olympic, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
  Root: camaro, Root Metric: 10, Metric: 20
    track-item: olympic.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
  Root: glacier, Root Metric: 10, Metric: 20
    track-item: olympic.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
camaro.00
Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
  Root: camaro, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
  Root: glacier, Root Metric: 10, Metric: 20
    track-item: camaro.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
  Root: olympic, Root Metric: 10, Metric: 20
    track-item: camaro.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
glacier.00
Primary next-hop: as0.0, IPV4, glacier
Primary next-hop: as0.0, IPV6, glacier
  Root: glacier, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
  Root: camaro, Root Metric: 10, Metric: 20
    track-item: glacier.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
  Root: olympic, Root Metric: 10, Metric: 20
    track-item: glacier.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
3 nodes

```



## show isis database

**List of Syntax**   [Syntax on page 235](#)  
                               [Syntax \(EX Series Switches and QFX Series\) on page 235](#)

**Syntax**   `show isis database`  
                   `<system-id>`  
                   `<brief | detail | extensive>`  
                   `<instance instance-name>`  
                   `<level (1 | 2)>`  
                   `<logical-system (all | logical-system-name)>`

**Syntax (EX Series Switches and QFX Series)**   `show isis database`  
   `<system-id>`  
   `<brief | detail | extensive>`  
   `<level (1 | 2)>`  
   `<instance instance-name>`

**Release Information**   Command introduced before Junos OS Release 7.4.  
                                   Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                   Command introduced in Junos OS Release 12.1 for the QFX Series.  
                                   Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  
                                   Command introduced in Junos OS Release 15.1X53-D30 for the QFX10002 switch.

**Description**   Display the entries in the Intermediate System-to-Intermediate System (IS-IS) link-state database, which contains data about PDU packets.

**Options**   **none**—Display standard information about IS-IS link-state database entries for all routing instances.

***system id***—(Optional) Display IS-IS link-state database entries for the specified intermediate system.

**brief | detail | extensive**—(Optional) Display the specified level of output.

**instance *instance-name***—(Optional) Display IS-IS link-state database entries for the specified routing instance.

**level (1 | 2)**—(Optional) Display IS-IS link-state database entries for the specified IS-IS level.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**Required Privilege Level**   view

**Related Documentation**   • [clear isis database on page 214](#)



**List of Sample Output** [show isis database on page 237](#)  
[show isis database brief on page 238](#)  
[show isis database detail on page 238](#)  
[show isis database extensive on page 239](#)  
[show isis database extensive \(SPRING\) on page 248](#)

**Output Fields** Table 11 on page 236 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 11: show isis database Output Fields*

| Field Name            | Field Description  | Level of Output         |
|-----------------------|--|-------------------------|
| <b>Interface name</b> | Name of the interface on which the link-state PDU has been received; always <b>IS-IS</b> for this command.   | All levels              |
| level                 | Level of intermediate system: <ul style="list-style-type: none"> <li>1—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system.</li> <li>2—Intermediate system routes between areas and toward other ASs.</li> </ul> | 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: <b>L1</b> , <b>L2</b> , <b>Overload</b> , or <b>Attached</b> (L1 only).  | none <b>brief</b>       |
| # LSPs                | Total number of link-state PDUs in the specified link-state database.  | none <b>brief</b>       |
| IP prefix             | Prefix advertised by this link-state PDU.  | <b>detail extensive</b> |
| IS neighbor           | IS-IS neighbor of the advertising system.  | <b>detail extensive</b> |
| IP prefix             | IPv4 prefix advertised by this link-state PDU.   | <b>detail extensive</b> |
| V6 prefix             | IPv6 prefix advertised by this link-state PDU.   | <b>detail extensive</b> |
| Metric                | Metric of the prefix or neighbor.  | <b>detail extensive</b> |



Table 11: show isis database Output Fields (continued)

| Field Name                   | Field Description  | Level of Output |
|------------------------------|--|-----------------|
| Header                       | <ul style="list-style-type: none"> <li>• <b>LSP ID</b>—Link state PDU identifier of the header.</li> <li>• <b>Length</b>—Header length.</li> <li>• <b>Allocated Length</b>—Amount of length available for the header.</li> <li>• <b>Router ID</b>—Address of the local routing device.</li> <li>• <b>Remaining Lifetime</b>—Remaining lifetime of the link-state PDU, in seconds.</li> </ul>   | extensive       |
| Packet                       | <ul style="list-style-type: none"> <li>• <b>LSP ID</b>—The identifier for the link-state PDU.</li> <li>• <b>Length</b>—Packet length.</li> <li>• <b>Lifetime</b>—Remaining lifetime, in seconds.</li> <li>• <b>Checksum</b>—The checksum of the link-state PDU.</li> <li>• <b>Sequence</b>—The sequence number of the link-state PDU. Every time the link-state PDU is updated, this number increments.</li> <li>• <b>Attributes</b>—Packet attributes.</li> <li>• <b>NLPID</b>—Network layer protocol identifier.</li> <li>• <b>Fixed length</b>—Specifies the set length for the packet.</li> </ul>  | extensive       |
| TLVs                         | <ul style="list-style-type: none"> <li>• <b>Area Address</b>—Area addresses that the routing device can reach.</li> <li>• <b>Speaks</b>—Supported routing protocols.</li> <li>• <b>IP router id</b>—ID of the routing device (usually the IP address).</li> <li>• <b>IP address</b>—IPv4 address.</li> <li>• <b>Hostname</b>—Assigned name of the routing device.</li> <li>• <b>IP prefix</b>—IP prefix of the routing device.</li> <li>• <b>Metric</b>—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device.</li> <li>• <b>IP extended prefix</b>—Extended IP prefix of the routing device.</li> <li>• <b>IS neighbor</b>—Directly attached neighbor's name and metric.</li> <li>• <b>IS extended neighbor</b>—Directly attached neighbor's name, metric, IP address, local interface index, and remote interface index.<br/><br/>The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.</li> <li>• <b>Router Capability</b>—ID of the routing device and flag.</li> </ul> <p><b>NOTE:</b> Router capability also specifies SPRING capability and SPRING algorithm when segment routing is enabled on the routing device.</p> | extensive       |
| Extended IS Reachability TLV | Type and length of the TLV is useful to identify the IS extended neighbors packed in this TLV.   | extensive       |
| SubTLV len                   | The length of the SubTLV of IS extended neighbor.  | 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 237](#).

### 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:
```

```

sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 970 secs
IS neighbor: hemantha-CE3.02 Metric: 10
Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
ES neighbor: 0015.0015.0015 Metric: 10 Down
ES neighbor: 0025.0025.0025 Metric: 10 Down
ES neighbor: 0030.0030.0030 Metric: 10 Down
ES neighbor: 0040.0040.0040 Metric: 10 Down
ES neighbor: sisira Metric: 0
IP prefix: 1.0.0.0/24 Metric: 10 External Down
IP prefix: 3.0.0.0/24 Metric: 10 External Down
IP prefix: 4.0.0.0/24 Metric: 10 External Down
IP prefix: 5.0.0.0/24 Metric: 10 Internal Up
IP prefix: 15.15.15.15/32 Metric: 10 External Down
IP prefix: 25.25.25.25/32 Metric: 10 External Down
IP prefix: 30.30.30.30/32 Metric: 10 External Down
IP prefix: 40.40.40.40/32 Metric: 10 External Down
IP prefix: 60.60.60.60/32 Metric: 0 Internal Up

```

```

Header: LSP ID: sisira.00-00, Length: 336 bytes
Allocated length: 336 bytes, Router ID: 0.0.0.0
Remaining lifetime: 970 secs, Level: 1, Interface: 333
Estimated free bytes: 144, Actual free bytes: 0
Aging timer expires in: 970 secs
Protocols: IP, IPv6, CLNS

```

```

Packet: LSP ID: sisira.00-00, Length: 336 bytes, Lifetime : 1198 secs
Checksum: 0x10fc, Sequence: 0x11, Attributes: 0xb L1 L2 Attached
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 18, Packet version: 1, Max area: 0

```

```
TLVs:
```

```

Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
Speaks: IP
Speaks: IPV6
Speaks: CLNP
Hostname: sisira
ES neighbor TLV: Internal, Metric: default 0, Up
ES: sisira
IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
IS extended neighbor: hemantha-CE3.02, Metric: default 10
ES neighbor TLV: External, Metric: default 10, Down
ES: 0040.0040.0040
ES neighbor TLV: External, Metric: default 10, Down

```



```

    ES: 0025.0025.0025
    ES neighbor TLV: External, Metric: default 10, Down
    ES: 0015.0015.0015
    ES neighbor TLV: External, Metric: default 10, Down
    ES: 0030.0030.0030
    IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
    IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
    IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
    IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
    IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
    IP external prefix: 1.0.0.0/24, Internal, Metric: default 10, Down
    IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down
    IP extended prefix: 3.0.0.0/24 metric 10 down
    IP extended prefix: 40.40.40.40/32 metric 10 down
    IP extended prefix: 4.0.0.0/24 metric 10 down
    IP extended prefix: 25.25.25.25/32 metric 10 down
    IP extended prefix: 15.15.15.15/32 metric 10 down
    IP extended prefix: 1.0.0.0/24 metric 10 down
    IP extended prefix: 30.30.30.30/32 metric 10 down
    IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
    IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
    IP extended prefix: 60.60.60.60/32 metric 0 up
    IP extended prefix: 5.0.0.0/24 metric 10 up
    No queued transmissions

```

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

```

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-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 database extensive (SPRING)

```

user@host> show isis database extensive
Extended IS Reachability TLV, Type: 22, Length: 99
IS extended neighbor: r0.00, Metric: default 10 SubTLV len: 88
  IP address: 11.1.1.2
  Neighbor's IP address: 11.1.1.1
  Local interface index: 333, Remote interface index: 342
  Current reservable bandwidth:
    Priority 0 : 1000Mbps
    Priority 1 : 1000Mbps
    Priority 2 : 1000Mbps
    Priority 3 : 1000Mbps
    Priority 4 : 1000Mbps
    Priority 5 : 1000Mbps
    Priority 6 : 1000Mbps
    Priority 7 : 1000Mbps
Maximum reservable bandwidth: 1000Mbps
Maximum bandwidth: 1000Mbps
Administrative groups: 0 <none>
P2P IPV6 Adj-SID - Flags:0xf0(F:1,B:1,V:1,L:1,S:0,P:0), Weight:0, Label: 401
P2P IPV6 Adj-SID: 401, Weight: 0, Flags: FBVL--
P2P IPV4 Adj-SID - Flags:0x70(F:0,B:1,V:1,L:1,S:0,P:0), Weight:0, Label:
400
P2P IPV4 Adj-SID: 400, Weight: 0, Flags: -BVL--
Extended IS Reachability TLV, Type: 22, Length: 144
IS extended neighbor: r1.06, Metric: default 10 SubTLV len: 133
  IP address: 12.1.1.1
  Local interface index: 336, Remote interface index: 0
  Current reservable bandwidth:
    Priority 0 : 1000Mbps
    Priority 1 : 1000Mbps
    Priority 2 : 1000Mbps
    Priority 3 : 1000Mbps
    Priority 4 : 1000Mbps
    Priority 5 : 1000Mbps
    Priority 6 : 1000Mbps
    Priority 7 : 1000Mbps
Maximum reservable bandwidth: 1000Mbps

```



```
Maximum bandwidth: 1000Mbps
Administrative groups: 0 <none>
LAN IPV6 Adj-SID -, Flags:0xfc(F:1,B:1,V:1,L:1,S:1,P:1), Weight:1
    Neighbor:r2, Label:803125
    LAN IPv6 Adj-SID: 803125, Weight: 1, Neighbor: r2, Flags: FBVLSP
    LAN IPV4 Adj-SID -, Flags:0x3c(F:0,B:0,V:1,L:1,S:1,P:1), Weight:1
        Neighbor:r2, Label:803122
    LAN IPv4 Adj-SID: 803122, Weight: 1, Neighbor: r2, Flags: --VLSP
    LAN IPV4 Adj-SID -, Flags:0x7c(F:0,B:1,V:1,L:1,S:1,P:1), Weight:1
        Neighbor:r2, Label:803123
    LAN IPv4 Adj-SID: 803123, Weight: 1, Neighbor: r2, Flags: -BVLSP
    LAN IPV6 Adj-SID -, Flags:0xf0(F:1,B:1,V:1,L:1,S:0,P:0), Weight:0
        Neighbor:r2, Label:37
    LAN IPv6 Adj-SID: 37, Weight: 0, Neighbor: r2, Flags: FBVL--
    LAN IPV4 Adj-SID -, Flags:0x70(F:0,B:1,V:1,L:1,S:0,P:0), Weight:0
        Neighbor:r2, Label:36
    LAN IPv4 Adj-SID: 36, Weight: 0, Neighbor: r2, Flags: -BVL--
```



## show isis hostname

**List of Syntax** [Syntax on page 250](#)  
[Syntax \(EX Series Switches and QFX Series\) on page 250](#)

**Syntax** `show isis hostname`  
`<logical-system (all | logical-system-name)>`

**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.  
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX 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 | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**Required Privilege Level** view

**List of Sample Output** [show isis hostname on page 251](#)

**Output Fields** [Table 12 on page 250](#) describes the output fields for the `show isis hostname` command. Output fields are listed in the approximate order in which they appear.

*Table 12: show isis hostname Output Fields*

| Field Name | Field Description                         |
|------------|---|
| System Id  | System identifier mapped to the hostname. |
| Hostname   | Hostname mapped to the system identifier. |



Table 12: show isis hostname Output Fields (continued)

| Field Name | Field Description   |
|------------|---|
| Type       | Type of mapping between system identifier and hostname. <ul style="list-style-type: none"><li>• <b>Dynamic</b>—Hostname mapping determined as described in RFC 2763, <i>Dynamic Hostname Exchange Mechanism for IS-IS</i>.</li><li>• <b>Static</b>—Hostname mapping configured by user.</li></ul> |

## Sample Output

### show isis hostname

```
user@host> show isis hostname
IS-IS hostname database:
System Id      Hostname      Type
1921.6800.4201 isis1         Dynamic
1921.6800.4202 isis2         Static
1921.6800.4203 isis3         Dynamic
```



## show isis interface

**List of Syntax**    [Syntax on page 252](#)  
                          [Syntax \(EX Series Switches and QFX Series\) on page 252](#)

**Syntax**    show isis interface  
                  <brief | detail | extensive>  
                  <interface-name>  
                  <logical-system (all | *logical-system-name*)>

**Syntax (EX Series Switches and QFX Series)**    show isis interface  
                  <brief | detail | extensive>  
                  <interface-name>

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

**Description**    Display status information about Intermediate System-to-Intermediate System (IS-IS)-enabled interfaces.



**NOTE:** If the configured metric for an IS-IS level is above 63, and the **wide-metrics-only** statement is not configured, the **show isis interface detail** command and the **show isis interface extensive** command display 63 as the metric value for that level. Configure the **wide-metrics-only** statement to generate metric values greater than 63 on a per IS-IS level basis.

The **show isis interface** command displays the configured metric value for an IS-IS level irrespective of whether is configured or not.

**Options**    **none**—Display standard information about all IS-IS-enabled interfaces.

**brief | detail | extensive**—(Optional) Display the specified level of output.

**interface-name**—(Optional) Display information about the specified interface only.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**Required Privilege Level**    view

**Related Documentation**    • *Understanding Wide IS-IS Metrics for Traffic Engineering*  
                                  • *Example: Enabling Wide IS-IS Metrics for Traffic Engineering*



**List of Sample Output** [show isis interface on page 255](#)  
[show isis interface brief on page 255](#)  
[show isis interface detail on page 255](#)  
[show isis interface extensive on page 255](#)

**Output Fields** [Table 13 on page 253](#) describes the output fields for the **show isis interface** command. Output fields are listed in the approximate order in which they appear.

*Table 13: show isis interface Output Fields*

| Field Name               | Field Description   | Level of Output         |
|--------------------------|---|-------------------------|
| <b>interface-name</b>    | Name of the interface.  | <b>detail</b>           |
| <b>Designated router</b> | Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.  | <b>detail</b>           |
| <b>Index</b>             | Interface index assigned by the Junos OS kernel.  | <b>detail</b>           |
| <b>State</b>             | Internal implementation information.  | <b>detail</b>           |
| <b>Circuit id</b>        | Circuit identifier.<br><br><b>NOTE:</b> Each IS-IS interface is assigned a circuit ID value to identify the interface within the linkstate database. All interfaces (loopback, broadcast, and so on) and all point-to-point links share the locally significant value of 0x01, and this value is not incremented.   | <b>detail</b>           |
| <b>Circuit type</b>      | Circuit type: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul>   | <b>detail</b>           |
| <b>LSP interval</b>      | Interval between link-state PDUs sent from the interface.   | <b>detail</b>           |
| <b>CSNP interval</b>     | Interval between complete sequence number PDUs sent from the interface.   | <b>detail extensive</b> |
| <b>Sysid</b>             | System identifier.  | <b>detail</b>           |
| <b>Interface</b>         | Interface through which the adjacency is made.  | <b>none brief</b>       |
| <b>L or Level</b>        | Level: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul> <b>NOTE:</b> The default IS-IS level on loopback interfaces are always same as the IS-IS level configured on other IS-IS interfaces in a router. You can also configure IS-IS level on loopback interfaces per your requirement. | All levels              |
| <b>CirID</b>             | Circuit identifier.   | <b>none brief</b>       |



Table 13: show isis interface Output Fields (continued)

| Field Name                         | Field Description   | Level of Output         |
|------------------------------------|---|-------------------------|
| Level 1 DR                         | Level 1 designated intermediate system.   | none <b>brief</b>       |
| Level 2 DR                         | Level 2 designated intermediate system.   | none <b>brief</b>       |
| L1/L2 Metric                       | Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0.  | none <b>brief</b>       |
| Flood-group Area-ID                | Flood-group is configured on a specific IS-IS interface<br><br><b>NOTE:</b> Seen only when Flood-group is configured.   | <b>detail extensive</b> |
| Adjacency advertisement: Advertise | This routing device has signaled to advertise this interface to its neighbors in their label-switched paths (LSPs).   | <b>detail extensive</b> |
| Adjacency advertisement: Suppress  | This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs.   | <b>detail extensive</b> |
| Adjacencies                        | Number of adjacencies established on this interface.  | <b>detail</b>           |
| Priority                           | Priority value for this interface.  | <b>detail</b>           |
| Metric                             | Metric value for this interface.  | <b>detail</b>           |
| Hello(s) / Hello Interval          | Interface's hello interval.   | <b>detail extensive</b> |
| Hold(s) / Hold Time                | Interface's hold time.  | <b>detail extensive</b> |
| Designated Router                  | Router responsible for sending network link-state advertisements, which describe all the routing devices attached to the network.   | <b>detail</b>           |
| Hello padding                      | Type of hello padding: <ul style="list-style-type: none"> <li>• <b>Adaptive</b>—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.</li> <li>• <b>Loose</b>—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state.</li> <li>• <b>Strict</b>—Padding is performed on all interface types and for all adjacency states, and is continuous.</li> </ul> | <b>extensive</b>        |
| LDP sync state                     | Current LDP synchronization state: <b>in sync</b> , <b>in holddown</b> , or <b>not supported</b> .  | <b>extensive</b>        |
| reason                             | Reason for being in the LDP sync state.   | <b>extensive</b>        |



Table 13: show isis interface Output Fields (continued)

| Field Name             | Field Description   | Level of Output  |
|------------------------|---|------------------|
| <b>config holdtime</b> | Configured value of the hold timer.   | <b>extensive</b> |
| <b>remaining</b>       | If the state is not in sync and the hold time is not infinity, then this field displays the remaining hold time in seconds. | <b>extensive</b> |
| <b>IIH max size</b>    | Configured value of IS-IS hello packets   | <b>extensive</b> |

## 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
lo0.0              3   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 255](#).

### 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
  Adjacency advertisement: Advertise
  Protection Type: Node Link, No eligible Backup
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1             0       64     0 Passive
    2             0       64     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, IIH max size:
  1505
  Adjacency advertisement: Advertise

  Flood-group Area-ID: 49.0001

```



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

|   |   |
|---|---|
| <b>Syntax</b>                                     | <pre>show isis overview &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis overview &lt;instance <i>instance-name</i>&gt;</pre>   |
| <b>Release Information</b>                        | <p>Command introduced in Junos OS Release 8.5.</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> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>   |
| <b>Description</b>                                | Display IS-IS overview information.   |
| <b>Options</b>                                    | <p><b>none</b>—Display standard overview information about IS-IS for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display overview information for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                   | view  |
| <b>List of Sample Output</b>                      | <a href="#">show isis overview on page 259</a>  |
| <b>Output Fields</b>                              | Table 14 on page 257 lists the output fields for the <b>show isis overview</b> command. Output fields are listed in the approximate order in which they appear.   |

Table 14: show isis overview Output Fields

| Field Name         | Field Description  |
|--------------------|--|
| Hostname           | Name of the router.  |
| Sysid              | Part of the ISO address of the routing device.                     |
| Areaid             | The area number of the routing device.                             |
| Instance           | IS-IS routing instance.  |
| Router ID          | Router ID of the routing device.                                   |
| Adjacency holddown | Adjacency holddown capability: <b>enabled</b> or <b>disabled</b> . |



Table 14: show isis overview Output Fields (continued)

| Field Name                      | Field Description  |
|---------------------------------|--|
| Maximum Areas                   | Maximum number of IS-IS areas advertised by the routing device.  |
| LSP life time                   | Lifetime of the link-state PDU, in seconds.  |
| Filter low life time LSPs up to | LSPs with a lifetime lower than this value are filtered out.   |
| Attached bit evaluation         | Attached bit capability: <b>enabled</b> or <b>disabled</b> .   |
| 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.   |
| 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: <b>enabled</b> or <b>disabled</b> .  |
| 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: <b>enabled</b> or <b>disabled</b> .  |
| Restart                         | Graceful restart capability: <b>enabled</b> or <b>disabled</b> .   |
| Restart duration                | Time period for complete reacquisition of IS-IS neighbors.   |
| Helper mode                     | Graceful restart helper capability: <b>enabled</b> or <b>disabled</b> .  |
| Level                           | IS-IS level: <ul style="list-style-type: none"> <li>• 1—Level 1 information</li> <li>• 2—Level 2 information</li> </ul>  |
| IPv4 is enabled                 | IP Protocol version 4 capability is enabled.   |
| IPv6 is enabled                 | IP Protocol version 6 capability is enabled.   |
| Micro-loop avoidance            | Micro-loop avoidance is enabled. Generally adjacent nodes converge faster than neighboring nodes causing traffic to loop. A route convergence delay is configured to avoid such micro loops. |
| Internal route preference       | Preference value of internal routes.   |



Table 14: show isis overview Output Fields (continued)

| Field Name                    | Field Description   |
|-------------------------------|---|
| External route preference     | Preference value of external routes.  |
| Prefix export limit           | Number of prefixes allowed to be exported, as configured by the <code>prefix-export-limit</code> 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.   |
| Adjacency holddown is active  | IS-IS adjacencies come up one after another when adjacency holddown is enabled.                             |

## Sample Output

### show isis overview

```

user@host> show isis overview
Instance: master
Router ID: 10.255.107.183
Hostname: pro-bng3-a
Sysid: 0192.0168.0001
Areaid: 49.0002
Adjacency holddown: enabled
Maximum Areas: 3
LSP life time: 1200
Filter low life time LSPs up to: 300
Attached bit evaluation: enabled
SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
IPv4 is enabled, IPv6 is enabled
Micro-loop avoidance: Enabled
  Method: Route Convergence Delay, Route convergence delay: 5000 msec
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
  Adjacency holddown is active
Level 2
  Internal route preference: 18
  External route preference: 165
  Prefix export limit: 5, Prefix export count: 5
  Wide metrics are enabled
  Adjacency holddown is active

```

```

user@host> show isis overview logical-system R2

```



```
Instance: master
  Router ID: 192.168.0.2
  Hostname: pro-bng3-a-R2
  Sysid: 0192.0168.0002
  Areaid: 49.0002
  Adjacency holddown: enabled
  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
  Prefix export count: 0
  Wide metrics are enabled, Narrow metrics are enabled
Level 2
  Internal route preference: 18
  External route preference: 165
  Prefix export count: 0
  Wide metrics are enabled, Narrow metrics are enabled
```

user@host> **show isis overview logical-system R3**

```
Instance: master
  Router ID: 192.168.0.3
  Hostname: pro-bng3-a-R3
  Sysid: 0192.0168.0003
  Areaid: 49.0002
  Adjacency holddown: enabled
  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
  Prefix export count: 0
  Wide metrics are enabled, Narrow metrics are enabled
Level 2
  Internal route preference: 18
  External route preference: 165
  Prefix export count: 0
  Wide metrics are enabled, Narrow metrics are enabled
```



## show isis route

|   |  |
|---|--|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 261</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 261</a>   |
| <b>Syntax</b>                                     | <pre>show isis route &lt;destination&gt; [inet   inet6] [instance instance-name] [logical-system (all   logical-system-name)] [topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)]</pre>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis route &lt;destination&gt; [inet   inet6] [instance instance-name] [topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)]</pre>   |
| <b>Release Information</b>                        | <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> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>  |
| <b>Description</b>                                | Display the routes in the IS-IS routing table.   |
| <b>Options</b>                                    | <p><b>none</b>—Display all routes in the IS-IS routing table for all supported address families for all routing instances.</p> <p><b>destination</b>—(Optional) Destination address for the route.</p> <p><b>inet   inet6</b>—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.</p> <p><b>instance instance-name</b>—(Optional) Display routes for the specified routing instance only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)</b>—(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.</p> |
| <b>Required Privilege Level</b>                   | view   |
| <b>List of Sample Output</b>                      | <a href="#">show isis route logical-system on page 262</a><br><a href="#">show isis route floodgroup on page 262</a><br><a href="#">show isis route (CLNS) on page 263</a><br><a href="#">show isis route on page 263</a>  |



**Output Fields** Table 15 on page 262 describes the output fields for the **show isis route** command. Output fields are listed in the approximate order in which they appear.

*Table 15: show isis route Output Fields*

| Field Name             | Field Description  |
|------------------------|--|
| <b>Current version</b> | Number of the current version of the IS-IS routing table.  |
| <b>L1</b>              | Version of Level 1 SPF that was run.   |
| <b>L2</b>              | Version of Level 2 SPF that was run.   |
| <b>Prefix or Label</b> | Destination of the route.  |
| <b>L</b>               | IS-IS level: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul> |
| <b>Version</b>         | Version of SPF that generated the route.   |
| <b>Metric</b>          | Metric value associated with the route.  |
| <b>Type</b>            | Metric type: <b>int</b> (internal) or <b>ext</b> (external).   |
| <b>Interface</b>       | Interface to the next hop.   |
| <b>Via</b>             | System identifier of the next hop, displayed as a name if possible.  |
| <b>ISO Routes</b>      | ISO routing table entries.   |
| <b>snpa</b>            | 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 floodgroup

```

user@R2> show isis route floodgroup 49.0001
IS-IS routing table           Current version: L1: 14 L2: 27
IPv4/IPv6 Routes

```



```

-----
Prefix          L  Version  Metric Type Interface    NH  Via
  Backup Score
0.0.0.0/0       1     14      10 int  ge-0/0/8.0    IPV4 R1
81.3.3.3/32     1     14      10 int  ge-0/0/8.0    IPV4 R1
128.220.17.202/32 1     14      10 int  ge-0/0/8.0    IPV4 R1

```

### 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 |

## MPLS Routes

-----

| Label                       | L | Version | Metric | Type | Interface   | NH   | Via            |
|-----------------------------|---|---------|--------|------|-------------|------|----------------|
| 300032 /52                  | 2 | 38      | 0      | int  | lt-1/2/0.13 | MPLS | Direct forward |
| to 10.0.7.60(pro-bng3-c-E)  |   |         |        |      |             |      |                |
| 300048 /52                  | 1 | 27      | 0      | int  | lt-1/2/0.12 | MPLS | Direct forward |
| to 10.0.6.60(pro-bng3-c-E)  |   |         |        |      |             |      |                |
| 300064 /52                  | 1 | 27      | 0      | int  | lt-1/2/0.12 | MPLS | Direct forward |
| to 10.0.6.60(pro-bng3-c-E)  |   |         |        |      |             |      |                |
| 300080 /52                  | 2 | 38      | 0      | int  | lt-1/2/0.12 | MPLS | Direct forward |
| to 10.0.6.60(pro-bng3-c-E)  |   |         |        |      |             |      |                |
| 300096 /52                  | 2 | 38      | 0      | int  | lt-1/2/0.12 | MPLS | Direct forward |
| to 10.0.6.60(pro-bng3-c-E)  |   |         |        |      |             |      |                |
| 299920 /52                  | 1 | 27      | 0      | int  | lt-1/2/0.14 | MPLS | Direct forward |
| to 10.0.10.70(pro-bng3-c-F) |   |         |        |      |             |      |                |
| 299936 /52                  | 1 | 27      | 0      | int  | lt-1/2/0.14 | MPLS | Direct forward |
| to 10.0.10.70(pro-bng3-c-F) |   |         |        |      |             |      |                |
| 299952 /52                  | 2 | 38      | 0      | int  | lt-1/2/0.14 | MPLS | Direct forward |
| to 10.0.10.70(pro-bng3-c-F) |   |         |        |      |             |      |                |
| 299968 /52                  | 2 | 38      | 0      | int  | lt-1/2/0.14 | MPLS | Direct forward |
| to 10.0.10.70(pro-bng3-c-F) |   |         |        |      |             |      |                |
| 299984 /52                  | 1 | 27      | 0      | int  | lt-1/2/0.13 | MPLS | Direct forward |
| to 10.0.7.60(pro-bng3-c-E)  |   |         |        |      |             |      |                |
| 300000 /52                  | 1 | 27      | 0      | int  | lt-1/2/0.13 | MPLS | Direct forward |
| to 10.0.7.60(pro-bng3-c-E)  |   |         |        |      |             |      |                |
| 300016 /52                  | 2 | 38      | 0      | int  | lt-1/2/0.13 | MPLS | Direct forward |
| to 10.0.7.60(pro-bng3-c-E)  |   |         |        |      |             |      |                |



## show isis spf

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 265</a><br><a href="#">Syntax (EX Series Switches) on page 265</a>   |
| <b>Syntax</b>                      | <pre>show isis spf (brief   log   results) &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>  |
| <b>Syntax (EX Series Switches)</b> | <pre>show isis spf (brief   log   results) &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>  |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>  |
| <b>Description</b>                 | Display information about IS-IS shortest-path-first (SPF) calculations.   |
| <b>Options</b>                     | <p><b>brief</b>—Display an overview of SPF calculations.</p> <p><b>log</b>—Display the log of SPF calculations.</p> <p><b>results</b>—Display the results of SPF calculations.</p> <p><b>instance <i>instance instance-name</i></b>—(Optional) Display SPF calculations for the specified routing instance.</p> <p><b>level (1   2)</b>—(Optional) Display SPF calculations for the specified IS-IS level.</p> <p><b>log</b>—Display the log of SPF calculations.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>results</b>—Display the results of SPF calculations.</p> <p><b>topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)</b>—(Optional) Display SPF calculations for the specified topology only.</p> |
| <b>Required Privilege Level</b>    | view  |
| <b>List of Sample Output</b>       | <a href="#">show isis spf log on page 266</a><br><a href="#">show isis spf results logical-system on page 267</a><br><a href="#">show isis spf results (CLNS) on page 268</a>   |



**Output Fields** Table 16 on page 266 describes the output fields for the **show isis spf** command. Output fields are listed in the approximate order in which they appear.

*Table 16: show isis spf Output Fields*

| Field Name     | Field Description   |
|----------------|---|
| Node           | System ID of a node.  |
| Metric         | Metric to the node.   |
| 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
           20      ge-1/1/0.0     scat     0:90:69:a6:48:9d
           10      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      gr-0/2/0.0     h
           30      8009:3::a09:3200/126
skag.02    20      gr-0/2/0.0     h
           10      gr-0/2/0.0     h
h.00       10      gr-0/2/0.0     h
           20      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      fe-0/0/1.0     toothache 0:12:0:34:0:56
           10      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      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

---

|  |  |
|--|--|
| List of Syntax                             | <a href="#">Syntax on page 270</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 270</a>   |
| Syntax                                     | <code>show isis statistics</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>  |
| Syntax (EX Series Switches and QFX Series) | <code>show isis statistics</code><br><code>&lt;instance <i>instance-name</i>&gt;</code>  |
| Release Information                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.<br>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.   |
| Description                                | Display statistics about IS-IS traffic.  |
| Options                                    | <b>none</b> —Display IS-IS traffic statistics for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display statistics for the specified routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(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"><li>• <a href="#">clear isis statistics on page 218</a></li></ul>  |
| List of Sample Output                      | <a href="#">show isis statistics on page 272</a>   |
| Output Fields                              | <a href="#">Table 17 on page 271</a> describes the output fields for the <b>show isis statistics</b> command. Output fields are listed in the approximate order in which they appear.  |



Table 17: show isis statistics Output Fields

| Field Name                  | Field Description   |
|-----------------------------|---|
| PDU type                    | <p>PDU type:</p> <ul style="list-style-type: none"> <li>• <b>CSNP</b>—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.</li> <li>• <b>IIH</b>—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.</li> <li>• <b>LSP</b>—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area.</li> <li>• <b>PSNP</b>—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.</li> <li>• <b>Unknown</b>—The PDU type is unknown.</li> </ul> |
| 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
```