JUNIPER NETWORKS

Junos® OS

IS-IS Configuration Guide

Release
12.1

Published: 2012-03-13

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

Junos® OS IS-IS Configuration Guide
12.1

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

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

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

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

Supported Platforms

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

- T Series
- MX Series
- M Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

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

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

Convention	Description	Examples
Italic text like this	• Introduces important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; interface names; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Enclose optional keywords or variables.	**stub** <**default-metric** *metric*>**;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** (***string1*** **|** ***string2*** **|** ***string3***)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Enclose a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identify a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

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

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/

- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

CHAPTER 1

Introduction to IS-IS

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.

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

This section discusses the following topics:

IS-IS Terminology

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 *protocol data units* (*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.

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
```

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, 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 value (TLV) in IS-IS link-state protocol data units (LSPs). This enables ISs in the routing domain to learn about the ISO system identifier of a particular IS.

IS-IS Packets

IS-IS uses the following protocol data units (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.
- Link-state PDUs (LSPs)—Contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area.
- 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.
- Partial sequence number PDUs (PSNPs)—Multicast by a receiver when it detects that it is missing a link-state PDU; that is, 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.

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 LSP regeneration.

IS-IS Extensions to Support Traffic Engineering

To help provide traffic engineering and MPLS with information about network topology and loading, extensions have been added to the Junos OS implementation of IS-IS. Specifically, IS-IS supports new TLVs that specify link attributes. These TLVs are included in the IS-IS link-state PDUs. The link-attribute information is used to populate the traffic engineering database, which is used by the Constrained Shortest Path First (CSPF) algorithm to compute the paths that MPLS LSPs take. This path information is used by RSVP to set up LSPs and reserve bandwidth for them.

NOTE: **Whenever possible, use IS-IS IGP shortcuts instead of traffic engineering shortcuts.**

The traffic engineering extensions are defined in Internet draft draft-isis-traffic-traffic-02, *IS-IS Extensions for Traffic Engineering*.

IS-IS IGP Shortcuts

In IS-IS, you can configure shortcuts, which allow IS-IS to use an LSP as the next hop as if it were a subinterface from the ingress routing device to the egress routing device. The address specified on the **to** statement at the **[edit protocols mpls label-switched-path**

***lsp-path-name*]** hierarchy level must match the router ID of the egress routing device for the LSP to function as a direct link to the egress routing device and to be used as input to IS-IS SPF calculations. When used in this way, LSPs are no different than Asynchronous Transfer Mode (ATM) and Frame Relay virtual circuits (VCs), except that LSPs carry only IPv4 traffic.

IS-IS Extensions to Support Route Tagging

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

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

CHAPTER 2

Introduction to IS-IS Configuration Guidelines

Overview of 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

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 may 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 may 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 may 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 may 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

- Configuring BFD Authentication for IS-IS on page 35
- **bfd-liveness-detection on page 111** statement
- **authentication-key-chains** statement in the *Junos OS System Basics Configuration Guide*
- **show bfd session** command in the *Junos OS Routing Protocols and Policies Command Reference*
- Configuring BFD for IS-IS on page 27

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 LSP. 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 10.

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

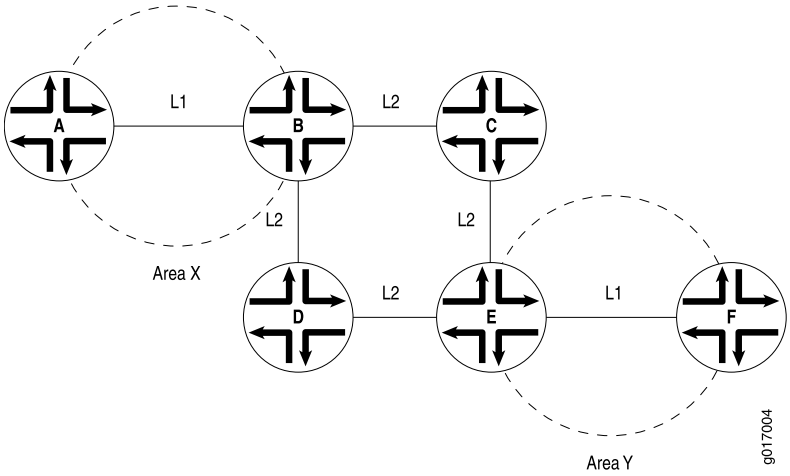

Overview of Hitless Authentication Key Rollover for IS-IS

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 21

PART 2

Configuration

CHAPTER 3

IS-IS Configuration Guidelines

Configuring IS-IS

To configure IS-IS, you include the following statements in the configuration:

```
protocols {
    isis {
        clns-routing;
        disable;
        ignore-attached-bit;
        graceful-restart {
            disable;
            helper-disable;
            restart-duration seconds;
        }
        label-switched-path name level level metric metric;
        level level-number {
            authentication-key key;
            authentication-key-chain key-chain-name;
            authentication-type authentication;
            external-preference preference;
            no-csnp-authentication;
            no-hello-authentication;
            no-psnp-authentication;
            preference preference;
            prefix-export-limit number;
            wide-metrics-only;
        }
        loose-authentication-check;
        lsp-lifetime seconds;
        max-areas seconds;
        no-adjacency-holddown;
```

```
no-authentication-check;
no-ipv4-routing;
no-ipv6-routing;
overload {
  advertise-high-metrics;
  timeout seconds;
}
reference-bandwidth reference-bandwidth;
rib-group {
  inet group-name;
  inet6 group-name;
}
spf-options {
  delay milliseconds;
  holddown milliseconds;
  rapid-runs number;
}
topologies {
  ipv4-multicast;
  ipv6-multicast;
  ipv6-unicast;
}
traffic-engineering {
  disable;
  ignore-lsp-metrics;
  family inet;
    shortcuts {
      multicast-rpf-routes;
    }
  }
  family inet6;
    shortcuts;
  }
}
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
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;
    multiplier number;
    no-adaptation;
    transmit-interval {
      threshold milliseconds;
      minimum-interval milliseconds;
```

```
            }
            version (1 | automatic);
          }
          checksum;
          csnp-interval (seconds | disable);
          hello-padding (adaptive | loose | strict);
          ldp-synchronization {
            disable;
            hold-time seconds;
          }
          link-protection;
          lsp-interval milliseconds;
          mesh-group (value | blocked);
          no-adjacency-holddown;
          no-eligible-backup;
          no-ipv4-multicast;
          no-ipv6-multicast;
          no-ipv6-unicast;
          no-unicast-topology;
          node-link-protection;
          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 number;
            ipv6-multicast-metric number;
            ipv6-unicast-metric number;
            metric metric;
            passive;
            priority number;
            te-metric metric;
          }
        }
      }
    }
```

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

By default, IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which an International Organization for Standardization (ISO) address is configured.

Minimum IS-IS Configuration

For IS-IS to run on the routing device, you must enable IS-IS on the routing device, configure a network entity title (NET) on one of the routing device's interfaces (preferably the loopback interface, **lo0**), and configure the ISO family on all interfaces on which you want IS-IS to run. When you enable IS-IS, Level 1 and Level 2 routes are enabled by default. The following is the minimum IS-IS configuration. In the **address** statement, ***address*** is the NET:

```
interfaces {
   lo0 {
      unit logical-unit-number {
         family iso {
            address address;
         }
      }
   }
   interface-type-fpc/pic/port {
      unit logical-unit-number {
         family iso;
      }
   }
}
protocols {
   isis {
      interface all;
   }
}
```

NOTE: To create the IS-IS interface, you must also configure IS-IS at the **[edit protocols isis interface *interface-name*]** hierarchy level. If you want Junos OS to create IS-IS interfaces automatically, include the **interface all** option at the **[edit protocols isis]** hierarchy level.

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 authentication for hello packets. To do this, see “Configuring Authentication for IS-IS Hello Packets” on page 58.

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 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) may 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 may 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 PSNP packets, include the **no-psnp-authentication** statement:

no-psnp-authentication;

To suppress authentication of CSNP packets, 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.

Example: Configuring Hitless Authentication Key Rollover for IS-IS

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

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 2 on page 22 shows the topology used in the example.

Figure 2: Hitless Authentication Key Rollover for IS-IS

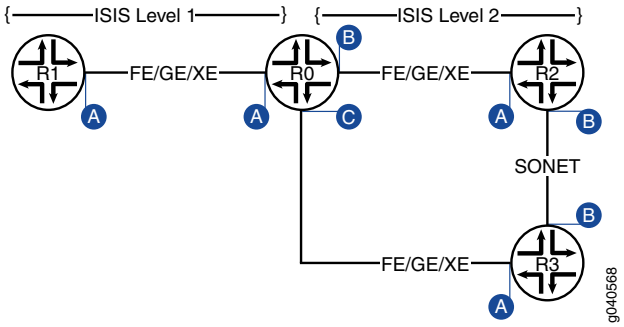

This example shows the configuration for Router R0.

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
  "$9$jfkqfTQnCpBDiCt"
set security authentication-key-chains key-chain base-key-global key 63 start-time
  "2011-8-6.06:54:00-0700"
set security authentication-key-chains key-chain base-key-global key 63 algorithm
  hmac-sha-1
set security authentication-key-chains key-chain base-key-global key 63 options
  isis-enhanced
set security authentication-key-chains key-chain base-key-inter key 0 secret
  "$9$8sgx7Vws4ZDkWLGD"
set security authentication-key-chains key-chain base-key-inter key 0 start-time
  "2011-8-6.06:54:00-0700"
set security authentication-key-chains key-chain base-key-inter key 0 algorithm md5
set security authentication-key-chains key-chain base-key-inter key 0 options basic
set protocols isis level 1 authentication-key-chain base-key-global
set protocols isis interface ge-0/0/0.0 level 1 hello-authentication-key-chain
  base-key-inter
```

Step-by-Step Procedure

To configure hitless authentication key rollover for IS-IS:

1. Configure the Router R0 interfaces.

```
[edit]
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 "$9$jfkqfTQnCpBDiCt"
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 "$9$8sgx7Vws4ZDkWLGD"
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 "$9$jfkqfTQnCpBDiCt"; ## SECRET-DATA
          start-time "2011-8-6.06:54:00-0700";
          algorithm hmac-sha-1;
          options isis-enhanced;
        }
      }
      key-chain base-key-inter {
        key 0 {
          secret "$9$8sgx7Vws4ZDkWLGD"; ## SECRET-DATA
          start-time "2011-8-6.06:54:00-0700";
          algorithm md5;
          options basic;
        }
      }
    }
```

Verification

To verify the configuration, run the following commands:

- **show isis authentication**
- **show security keychain**

Related Documentation

- Overview of Hitless Authentication Key Rollover for IS-IS on page 10

Configuring of Interface-Specific IS-IS Properties

You can configure interface-specific IS-IS properties by including the **interface** statement.

```
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;
    version (1 | automatic);
  }
  checksum;
  csnp-interval (seconds | disable);
  ldp-synchronization {
    disable;
    hold-time seconds;
  }
  lsp-interval milliseconds;
  mesh-group (value | blocked);
  no-ipv4-multicast;
  no-ipv6-multicast;
  no-ipv6-unicast;
  no-unicast-topology;
  passive;
  point-to-point;
  level level-number {
    disable;
    hello-authentication-type authentication;
    hello-authentication-key key;
    hello-authentication-key-chain key-chain-name;
    hello-interval seconds;
    hold-time seconds;
    ipv4-multicast-metric number;
    ipv6-multicast-metric number;
    ipv6-unicast-metric number;
    metric metric;
    passive;
    priority number;
    te-metric metric;
```

```
    }
}
```

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

For ***interface-name***, specify the full interface name, including the physical and logical address components. To configure all interfaces, specify the interface name as **all**.

For more information about configuring IS-IS interface properties, see the following topics:

- Configuring BFD for IS-IS on page 27
- Overview of BFD Authentication for IS-IS on page 7
- Configuring BFD Authentication for IS-IS on page 35
- Enabling Packet Checksums on IS-IS Interfaces on page 38
- Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces on page 38
- Configuring Synchronization Between LDP and IS-IS on page 38
- Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces on page 39
- Configuring Mesh Groups of IS-IS Interfaces on page 39
- Configuring IS-IS Multicast Topology on page 40
- Configuring IS-IS IPv6 Unicast Topologies on page 55
- Configuring Point-to-Point Interfaces for IS-IS on page 56
- Configuring Levels on IS-IS Interfaces on page 56

Configuring BFD for IS-IS

- Overview of Configuring BFD for IS-IS on page 27
- Example: Configuring BFD for IS-IS on page 29

Overview of Configuring BFD for IS-IS

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

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15

seconds. A back-off algorithm increases the receive (RX) interval by two if the local BFD instance is the reason for the session flap. The transmission (TX) interval is increased by two if the remote BFD instance is the reason for the session flap.

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

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

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

Table 3: Configuring BFD for IS-IS

Statement	Description
bfd-liveness-detection	Enable failure detection.
minimum-interval *milliseconds*	Specify the minimum transmit and receive intervals for failure detection. This value represents the minimum interval at which the local router transmits hellos packets as well as the minimum interval at which the router expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately. NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping. Depending on your network environment, these additional recommendations might apply: • For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions. • For very large-scale network deployments with a large number of BFD sessions, please contact Juniper Networks customer support for more information. • For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.
minimum-receive-interval *milliseconds*	Specify only the minimum receive interval for failure detection. This value represents the minimum interval at which the local router expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number from 1 through 255,000 milliseconds.
multiplier *number*	Specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down. The default is 3, and you can configure a value from 1 through 225.

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

Statement	Description
no-adaptation	Disable BFD adaptation. In Junos OS Release 9.0 and later, you can specify that the BFD sessions not adapt to changing network conditions. NOTE: We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.
threshold	• Specify the threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent. • Specify the threshold for the transmit interval. NOTE: The threshold value must be greater than the minimum transmit interval multiplied by the multiplier number.
transmit-interval minimum-interval	Specify the minimum transmit interval for failure detection. This value represents the minimum interval at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value from 1 through 255,000 milliseconds.
version	Specify the BFD version used for detection. The default is to have the version detected automatically.

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

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

Example: Configuring BFD for IS-IS

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

Requirements

Before you begin, configure IS-IS on both routers. See "Minimum IS-IS Configuration" on page 18 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 3 on page 30 shows the sample network.

Figure 3: Configuring BFD on IS-IS

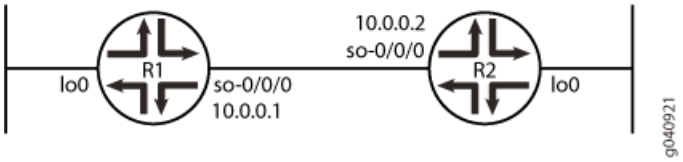

Configuration

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 the *Junos OS CLI User Guide*.

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 Confirm your configuration by issuing the **show protocols isis interface** command.

```
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

Purpose Make sure that the 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 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
                                                Detect   Transmit
Address                 State     Interface     Time     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
```



```
                                                  Detect   Transmit
Address                 State     Interface       Time     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.

Configuring BFD Authentication for IS-IS

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

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

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

- Configuring BFD Authentication Parameters on page 35
- Viewing Authentication Information for BFD Sessions on page 36

Configuring BFD Authentication Parameters

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on an IS-IS route or routing instance.

   ```
   [edit]
   user@host# set protocols isis interface if1-isis bfd-liveness-detection authentication
     algorithm keyed-sha-1
   ```

   

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

2. Specify the keychain to be used to associate BFD sessions on the specified IS-IS route or routing instance with the unique security authentication keychain attributes. This should match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit]
user@host# set protocols isis interface if1-isis bfd-liveness-detection authentication
  keychain bfd-isis
```

NOTE: The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:
   - The matching *key-chain-name* as specified in Step 2.
   - At least one *key*, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
   - The *secret-data* used to allow access to the session.
   - The time at which the authentication key becomes active, *yyyy-mm-dd.hh:mm:ss*.

   ```
   [edit security]
   user@host# authentication-key-chains key-chain bfd-sr4 key 53 secret
     $9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm start-time 2009-06-14.10:00:00
   ```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

   ```
   [edit]
   user@host> set protocols isis interface if1-isis bfd-liveness-detection authentication
     loose-check
   ```

5. (Optional) View your configuration using the **show bfd session detail** or **show bfd session extensive** command.
6. Repeat these steps to configure the other end of the BFD session.

NOTE: BFD authentication is only supported in the domestic image and is not available in the export image.

Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **if1-isis** interface. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-isis**. The authentication keychain is configured with two keys. Key **1** contains the secret data "**9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm**" and a start time of June 1, 2009, at 9:46:02 AM PST. Key **2** contains the secret data "**9a5jiKW9l.reP38ny.TszF2/9**" and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols isis]
interface if1-isis {
```

```
    bfd-liveness-detection {
      authentication {
        algorithm keyed-sha-1;
        key-chain bfd-isis;
      }
    }
  }
  [edit security]
  authentication key-chains {
    key-chain bfd-isis {
      key 1 {
        secret "$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm";
        start-time "2009-6-1.09:46:02 -0700";
      }
      key 2 {
        secret "$9$a5jiKW9l.reP38ny.TszF2/9";
        start-time "2009-6-1.15:29:20 -0700";
      }
    }
  }
```

If you commit these updates to your configuration, you see output similar to the following. In the output for the **show bfd sessions detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the **show bfd sessions extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

show bfd sessions detail

```
user@host# show bfd session detail

                                                Detect   Transmit
Address                 State     Interface     Time     Interval  Multiplier
10.9.1.29                Up        ge-4/0/0.0     0.600     0.200       3
 Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3, Authenticate

 Session up time 3d 00:34, previous down time 00:00:01
 Local diagnostic NbrSignal, remote diagnostic AdminDown
 Remote state Up, version 1

1 sessions, 1 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

show bfd sessions extensive

```
user@host# show bfd session extensive
                                                Detect   Transmit
Address                 State     Interface     Time     Interval  Multiplier
10.9.1.29            Up        ge-4/0/0.0          0.600     0.200      3
  Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3, Authenticate

          keychain bfd-isis, algo keyed-sha-1, mode strict
  Session up time 00:04:42
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated
 Min async interval 0.300, min slow interval 1.000
 Adaptive async TX interval 0.300, RX interval 0.300
 Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3
 Remote min TX interval 0.300, min RX interval 0.300, multiplier 3
```

```
Local discriminator 2, remote discriminator 2
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-isis, algo keyed-sha-1, mode strict

1 sessions, 1 clients

Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

Related Documentation

- Overview of BFD Authentication for IS-IS on page 7
- bfd-liveness-detection on page 111
- **authentication-key-chains** statement in the *Junos OS System Basics Configuration Guide*
- **show bfd session** command in the *Junos OS Routing Protocols and Policies Command Reference*
- Configuring BFD for IS-IS on page 27

Enabling Packet Checksums on IS-IS Interfaces

You can enable checksums for packets on a per-interface basis. To enable checksums, include the **checksum** statement:

```
checksum;
```

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

Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces

By default, IS-IS sends complete sequence number (CSN) packets periodically. If the routing device is the designated router on a LAN, IS-IS sends CSN packets every 10 seconds. If the routing device is on a point-to-point interface, it sends CSN packets every 5 seconds. You might want to modify the default interval to protect against link-state PDU (LSP) flooding.

To modify the CSNP interval, include the **csnp-interval** statement:

```
csnp-interval seconds;
```

The time can range from 1 through 65,535 seconds.

To configure the interface not to send any CSN packets, specify the **disable** option:

```
csnp-interval disable;
```

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

Configuring Synchronization Between LDP and IS-IS

LDP distributes labels in non-traffic-engineered applications. Labels are distributed along the best path determined by IS-IS. If the synchronization between LDP and IS-IS is lost, the label-switched path (LSP) goes down. Therefore, LDP and IS-IS synchronization is

beneficial. When LDP is not fully operational on a given link (a session is not established and labels are not exchanged), IS-IS advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.

LDP synchronization is supported only on point-to-point interfaces and LAN interfaces configured as point-to-point interfaces under IS-IS. LDP synchronization is not supported during graceful restart.

To advertise the maximum cost metric until LDP is operational for LDP synchronization, include the **ldp-synchronization** statement:

```
ldp-synchronization {
  disable;
  hold-time seconds;
}
```

To disable synchronization, include the **disable** statement. To configure the time period to advertise the maximum cost metric for a link that is not fully operational, include the **hold-time** statement.

NOTE: When an interface has been in the holddown state for more than 3 minutes, a system log message with a warning level is sent. This message appears in both the messages file and the trace file.

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

Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces

By default, the routing device sends one link-state PDU packet out an interface every 100 milliseconds. To modify this interval, include the **lsp-interval** statement:

```
lsp-interval milliseconds;
```

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

To disable the transmission of all link-state PDU packets, set the interval to 0.

Configuring Mesh Groups of IS-IS Interfaces

A *mesh group* is a set of routing devices that are fully connected; that is, they have a fully meshed topology. When link-state PDU packets are being flooded throughout an area, each router within a mesh group receives only a single copy of a link-state PDU packet instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of link-state PDU packets.

To create a mesh group and designate that an interface be part of the group, assign a mesh-group number to all the routing device interfaces in the group:

```
mesh-group value;
```

To prevent an interface in the mesh group from flooding link-state PDUs, configure blocking on that interface:

```
mesh-group blocked;
```

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

Configuring IS-IS Multicast Topology

IS-IS 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 lets you exercise control over the paths that multicast data takes through a network so that it is independent of unicast data paths. You can also configure IS-IS to calculate an alternate IPv6 multicast topology, in addition to the normal IPv6 unicast topology.

NOTE: **IS-IS only starts advertising the routes when the interface routes are in inet.2.**

Table 4 on page 40 lists the various IPv4 statements you can use to configure IS-IS topologies.

Table 4: IPv4 Statements

Statement	Description
ipv4-multicast	Enables an alternate IPv4 multicast topology.
ipv4-multicast-metric ***number***	Configures the multicast metric for an alternate IPv4 multicast topology.

Table 4: IPv4 Statements *(continued)*

Statement	Description
no-ipv4-multicast	Excludes an interface from the IPv4 multicast topology.
no-unicast-topology	Excludes an interface from the IPv4 unicast topologies.

Table 5 on page 41 lists the various IPv6 statements you can use to configure IS-IS topologies.

Table 5: IPv6 Statements

Statement	Description
ipv6-multicast	Enables an alternate IPv6 multicast topology.
ipv6-multicast-metric ***number***	Configures the multicast metric for an alternate IPv6 multicast topology.
ipv6-unicast-metric ***number***	Configures the unicast metric for an alternate IPv6 multicast topology.
no-ipv6-multicast	Excludes an interface from the IPv6 multicast topology.
no-ipv6-unicast	Excludes an interface from the IPv6 unicast topologies.

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

Example: Configuring IS-IS Multicast Topology

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

Requirements

Before you begin, configure IS-IS on all routers. See "Minimum IS-IS Configuration" on page 18 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 4 on page 42 shows the sample network.

Figure 4: Configuring IS-IS Multicast Topology

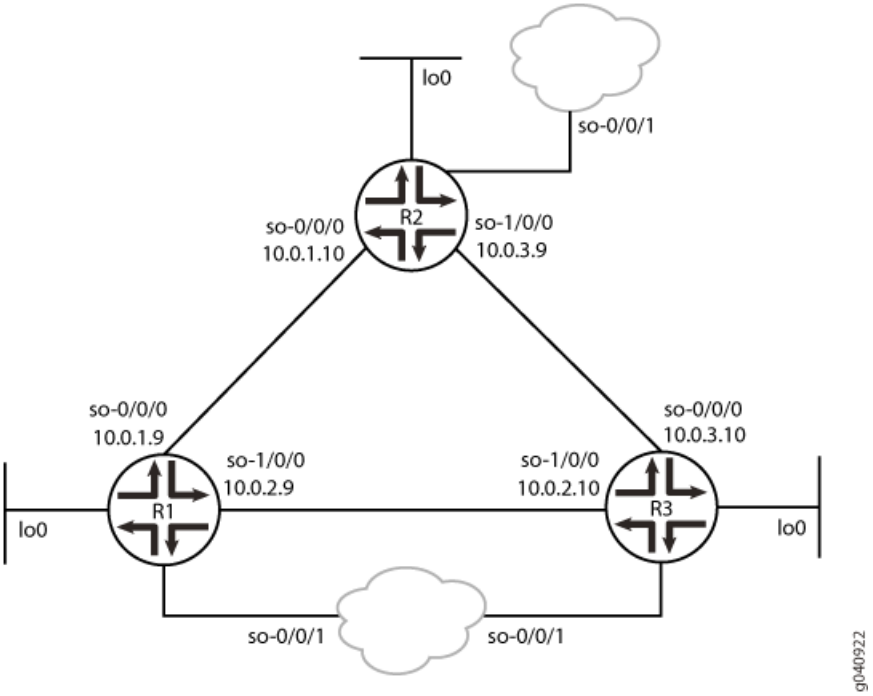

Configuration

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 the *Junos OS 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 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 19
   ```

```
user@R2# set level 1 ipv4-multicast-metric 11
user@R2# set level 2 metric 27
user@R2# set level 2 ipv4-multicast-metric 21
```

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 Confirm your configuration by using the **show protocols isis** statement.

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 46
- Verifying That IS-IS Is Configured on page 48
- Verifying the Configured Multicast Metric Values on page 50
- Verifying the Configuration of the Multicast Topology on page 51

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.

Configuring IS-IS 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 lets you exercise control over the paths that unicast data takes through a network.

To enable an alternate IPv6 unicast topology for IS-IS, include the **ipv6-unicast** statement:

```
isis {
  topologies {
    ipv6-unicast;
  }
}
```

To configure a metric for an alternate 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;
  }
}
```

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

Configuring Point-to-Point Interfaces for IS-IS

You can use the **point-to-point** statement to configure a LAN interface to act like a point-to-point interface for IS-IS. You do not need an unnumbered LAN interface, and it has no effect if configured on an interface that is already point-to-point.

The **point-to-point** statement affects only IS-IS protocol procedures on that interface; all other protocols continue to treat the interface as a LAN interface. Only two IS-IS routing devices can be connected to the LAN interface and both must be configured as point-to-point.

To configure a point-to-point IS-IS interface, include the **point-to-point** statement:

```
point-to-point;
```

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

Configuring Levels on IS-IS Interfaces

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. In order to route between areas you must have two adjacent Level 2 routers that communicate with each other.

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.

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.

To configure an area, include the **level** statement:

```
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 number;
    ipv6-multicast-metric number;
    ipv6-unicast-metric number;
    metric metric;
    passive;
    priority number;
    te-metric metric;
}
```

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

The statements within the **level** statement allow you to perform the following tasks when you are configuring the following optional level-specific properties:

- Disabling IS-IS at a Level on IS-IS Interfaces on page 57
- Advertising Interface Addresses Without Running IS-IS on page 58
- Configuring Authentication for IS-IS Hello Packets on page 58
- Configuring the Transmission Frequency for IS-IS Hello Packets on page 59
- Configuring the Delay Before IS-IS Neighbors Mark the Routing Device as Down on page 59
- Configuring the Metric Value for IS-IS Routes on page 60
- Configuring the IS-IS Metric Value Used for Traffic Engineering on page 60
- Configuring the Designated Router Priority for IS-IS on page 61

Disabling IS-IS at a Level on IS-IS Interfaces

By default, IS-IS is enabled for IS-IS areas on all enabled interfaces on which the ISO protocol family is enabled (at the **[edit interfaces** *interface* **unit** *logical-unit-number***]** hierarchy level). To disable IS-IS at any particular level on an interface, include the **disable** statement:

```
disable;
```

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

Enabling IS-IS on an interface (by including the **interface** statement at the **[edit protocols isis]** hierarchy level), disabling it (by including the **disable** statement), and not actually having IS-IS run on an interface (by including the **passive** statement) are mutually exclusive states.

Example: Disabling IS-IS at a Level

On SONET/SDH interface **so-0/0/0**, enable IS-IS for Level 1 only. With this configuration, tracing messages periodically indicate that IS-IS is creating Level 2 link-state PDUs. However, because IS-IS for Level 2 is disabled, these link-state PDUs are never distributed to neighboring routers.

```
protocols {
```

```
isis {
  traceoptions {
    file isis size 1m files 10;
    flag spf;
    flag lsp;
    flag error;
  }
  interface so-0/0/0 {
    level 2 {
      disable;
    }
  }
}
```

Advertising Interface Addresses Without Running IS-IS

By default, IS-IS must be configured on an interface or a level for direct interface addresses to be advertised into that level. To advertise the direct interface addresses without actually running IS-IS on that interface or level, include the **passive** statement:

```
passive;
```

NOTE: Configuring IS-IS on a loopback interface automatically renders it as a passive interface, irrespective of whether the passive statement was used in the configuration of the interface.

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

Enabling IS-IS on an interface (by including the **interface** statement at the **[edit protocols isis]** hierarchy level), disabling it (by including the **interface disable** statement), and not actually having IS-IS run on an interface (by including the **passive** statement) are mutually exclusive states.

NOTE: If neither passive mode nor the family iso option is configured on the IS-IS interface, then the routing device treats the interface as not being operational, and no direct IPv4/IPv6 routes are exported into IS-IS. (You configure the family iso option at the **[edit interfaces** ***interface-name*** **unit** ***logical-unit-number***] hierarchy level.)

Configuring Authentication for IS-IS Hello Packets

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.

CAUTION: If no authentication is configured for Level 1 on a point-to-point link with both levels enabled, the hello packets are sent without any password, regardless of the Level 2 authentication configurations.

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.

To enable hello authentication for an IS-IS level on an interface and define the password, include the **hello-authentication-type** and **hello-authentication-key** statements. To configure hitless authentication key rollover, include the **hello-authentication-key-chain** statement:

```
hello-authentication-type (md5 | simple);
hello-authentication-key password;
hello-authentication-key-chain key-chain-name;
```

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

Configuring the Transmission Frequency for IS-IS Hello Packets

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. By default, a designated intermediate system (DIS) router sends hello packets every 3 seconds, and a non-DIS router sends hello packets every 9 seconds.

To modify how often the routing device sends hello packets out of an interface, include the **hello-interval** statement:

```
hello-interval seconds;
```

The hello interval range is from 1 through 20,000 seconds.

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

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

Configuring the Delay Before IS-IS Neighbors Mark the Routing Device as Down

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. The default hold-time value is three times the default hello interval: 9 seconds for a designated intermediate system (DIS) router and 27 seconds for a non-DIS router.

To modify the hold-time value on the local routing device, include the **hold-time** statement:

```
hold-time seconds;
```

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

Configuring the Metric Value for IS-IS Routes

All IS-IS routes have a cost, which is a routing metric that is used in the IS-IS link-state calculation. The cost is an arbitrary, dimensionless integer that can be from 1 through 63, or from 1 through 16,777,215 ($2^{24} - 1$) if you are using wide metrics. The default metric value is 10 (with the exception of the **lo0** interface, which has a default metric of 0). To modify the default value, include the **metric** statement:

metric *metric*;

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

Table 6 on page 60 depicts IS-IS route export metric default values.

Table 6: 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

NOTE: The default metric values behavior can be customized by using the Routing Policy framework.

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

For more information about IS-IS interface metric values, see "Configuring the Reference Bandwidth Used in IS-IS Metric Calculations" on page 61.

Configuring the IS-IS Metric Value Used for Traffic Engineering

When traffic engineering is enabled on the routing device, you can configure an IS-IS metric that is used exclusively for traffic engineering. The traffic engineering metric is

used for information injected into the traffic engineering database. Its value does not affect normal IS-IS forwarding.

To modify the default value, include the **te-metric** statement:

```
te-metric metric;
```

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

Configuring the Designated Router Priority for IS-IS

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.

The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.

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. By default, routing devices have a priority value of 64.

To modify the interface's priority value, include the **priority** statement:

```
priority number;
```

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

Configuring the Reference Bandwidth Used in IS-IS Metric Calculations

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.

The cost of a route is described by a single dimensionless metric that is determined using the following formula:

```
cost = reference-bandwidth/bandwidth
```

To modify the reference bandwidth, include the **reference-bandwidth** statement:

```
reference-bandwidth reference-bandwidth;
```

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

reference-bandwidth is the reference bandwidth. If the reference bandwidth is not configured, all interfaces have a default metric of 10 (with the exception of the **lo0** interface, which has a default metric of 0).

For example, if you set the reference bandwidth to 1 Gbps (that is, ***reference-bandwidth*** is set to 1,000,000,000), a 100-Mbps interface has a default metric of 10.

For more information about IS-IS route metrics, see "Configuring the Metric Value for IS-IS Routes" on page 60.

Limiting the Number of Advertised IS-IS Areas

By default, IS-IS advertises a maximum of three areas in the IS-IS hello (IIH) PDUs and link-state PDUs. To advertise more than three ISO network addresses for a router, include the **max-areas** statement:

```
max-areas number;
```

The range that you can configure is from 3 through 36, and the default is 3. This value is included in the Maximum Address Area field of the IS-IS common PDU header included in all outgoing PDUs.

NOTE: 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.

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

Enabling Wide IS-IS Metrics for Traffic Engineering

Normally, IS-IS metrics can have values up to 63, and IS-IS generates two type length values (TLVs), one for an IS-IS adjacency and the second for an IP prefix. To allow IS-IS to support traffic engineering, a second pair of TLVs has been added to IS-IS, one for IP prefixes and the second for IS-IS adjacency and traffic engineering information. With these TLVs, IS-IS metrics can have values up to 16,777,215 ($2^{24} - 1$).

By default, Junos OS supports the sending and receiving of wide metrics. Junos OS allows a maximum metric value of 63 and generates both pairs of TLVs. To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, include the **wide-metrics-only** statement:

```
wide-metrics-only;
```

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

Configuring Preference Values for IS-IS Routes

Route preferences 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. For more information about route preferences, see Route Preferences Overview.

By default, Level 1 IS-IS internal routes have a preference value of 15, Level 2 IS-IS internal routes have a preference of 18, Level 1 IS-IS external routes have a preference of 160, and Level 2 external routes have a preference of 165. To change the preference values, include the **preference** statement (for internal routes) or the **external-preference** statement:

```
external-preference preference;
preference preference;
```

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

The preference value can range from 0 through 4,294,967,295 ($2^{32} - 1$).

Limiting the Number of Prefixes Exported to IS-IS

By default, there is no limit to the number of prefixes that can be exported into IS-IS. To configure a limit to the number of prefixes that can be exported into IS-IS, include the **prefix-export-limit** statement:

```
prefix-export-limit number;
```

If the number of prefixes exported into IS-IS exceeds the configured limit, the overload bit is set and the overload state is reached. The overload state can then be cleared by using the **clear isis overload** command.

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

You can specify a number range from 0 through 4,294,967,295.

Configuring the Link-State PDU Lifetime for IS-IS

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 *LSP lifetime*, normally is sufficient to guarantee that link-state PDUs never expire.

To modify the link-state PDU lifetime, include the **lsp-lifetime** statement:

```
lsp-lifetime seconds;
```

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

The time can range from 350 through 65,535 seconds.

The link-state PDU refresh interval is derived from the link-state PDU lifetime and is equal to the lifetime minus 317 seconds.

Advertising Label-Switched Paths into IS-IS

You can advertise label-switched paths into IS-IS as point-to-point links, and the label-switched paths can be used in SPF calculations. The advertisement contains a

local address (the **from** address of the label-switched path), a remote address (the **to** address of the label-switched path), and a metric with the following precedence:

- Use the label-switched path metric defined under IS-IS.
- Use the label-switched path metric configured for the label-switched path under MPLS.
- If you do not configure any of the above, use the default IS-IS metric of 10.

To advertise label-switched paths, include the **label-switched-path** statement, with a specified **level** and **metric**:

label-switched-path *name* level *level* metric *metric*;

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

NOTE: **Before a single-hop label-switched path between a multiaccess link can be announced as up and used in SPF calculations, you must configure a label-switched path in both directions between two label-switched routers.**

Related Documentation

- *Junos OS MPLS Applications Configuration Guide*

Configuring IS-IS to Make Routing Devices Appear Overloaded

If the time elapsed after the IS-IS instance is enabled is less than the specified timeout, overload mode is set.

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.

A timer is started for the difference between the timeout and the time elapsed since the instance started. 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 avoid the overloaded routing device and take paths around the routing device. However, the overloaded routing device's own links are still accessible.

To summarize, the value of the overload bit depends on these three scenarios:

1. When the overload bit has already been set to a given value and the routing process is restarted: LSPs are regenerated with the overload bit cleared.
2. When the overload bit is reset to a lesser value while the routing process is running: LSPs are regenerated with the overload bit cleared.
3. When the overload bit is reset to a greater value while the routing process is running: LSPs are regenerated with the overload bit set to the difference between the old and new value.

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. However, the overloaded routing device's own links are still accessible.

You can configure the local routing device so that it appears to be overloaded. You might want to do this when you want the routing device to participate in IS-IS routing, but do not want it to be used for transit traffic. (Note that traffic to immediately attached interfaces continues to transit the routing device.) To mark the routing device as overloaded, include the **overload** statement:

```
overload {
    advertise-high-metrics;
    allow-route-leaking;
    timeout seconds;
}
```

To advertise maximum link metrics in network layer reachability information (NLRI) instead of setting the overload bit, include the **advertise-high-metrics** option when specifying the **overload** statement:

```
advertise-high-metrics;
```

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

To allow route information to pass (leak) into the network when the routing device is in overload mode, include the **allow-route-leaking** option when specifying the **overload** statement:

```
allow-route-leaking;
```

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.

To specify the number of seconds at which overload is reset, include the **timeout** option when specifying the **overload** statement:

```
overload timeout seconds;
```

The time can range from 60 through 1800 seconds.

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

Configuring SPF Options for IS-IS

You can configure the following shortest-path-first (SPF) options:

- The delay in the time between the detection of a topology change and when the SPF algorithm actually runs.
- The maximum number of times that the SPF algorithm can run in succession before the hold-down timer begins.
- The time to hold down, or wait, before running another SPF calculation after the SPF algorithm has run in succession the configured maximum number of times.

To configure SPF options, include the **spf-options** statement:

```
spf-options {
    delay milliseconds;
    holddown milliseconds;
    rapid-runs number;
}
```

To configure the SPF delay, include the **delay** statement when specifying the **spf-options** statement:

```
delay milliseconds;
```

By default, the SPF algorithm runs 200 milliseconds after the detection of a topology change. The range that you can configure is from 50 through 1000 milliseconds.

To configure the maximum number of times that the SPF algorithm can run in succession, include the **rapid-runs** statement when specifying the **spf-options** statement:

```
rapid-runs number;
```

The default number of SPF calculations that can occur in succession is 3. The range that you can configure is from 1 through 5. Each SPF algorithm is run after the configured SPF delay. When the maximum number of SPF calculations occurs, the hold-down timer begins. Any subsequent SPF calculation is not run until the hold-down timer expires.

To configure the SPF hold-down timer, include the **holddown** statement when specifying the **spf-options** statement:

```
holddown milliseconds;
```

The default is 5000 milliseconds, and the range that you can configure is from 2000 through 10,000 milliseconds. Use the hold-down timer to hold down, or wait, before running any subsequent SPF calculations after the SPF algorithm runs for the configured maximum number of times. If the network stabilizes during the hold-down period and the SPF algorithm does not need to run again, the system reverts to the configured values for the **delay** and **rapid-runs** statements.

Configuring Graceful Restart for IS-IS

Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled globally for all routing protocols at the **[edit routing-options]** hierarchy level. When graceful restart for IS-IS 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.

You can configure graceful restart parameters specifically for IS-IS. To do this, include the **graceful-restart** statement:

```
graceful-restart {
   helper-disable;
   restart-duration seconds;
}
```

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

To disable graceful restart for IS-IS, specify the **disable** statement. Helper mode is enabled by default. To disable the graceful restart helper capability, specify the **helper-disable** statement. To configure a time period for complete restart, specify the **restart-duration** statement. You can specify a number between 1 and 3600. The default value is 90 seconds.

On LAN interfaces where IS-IS is configured on a transit router that serves as the designated router (DR), a graceful restart causes:

- The ingress router of the label-switched path (LSP), which passes through the DR, to break the LSP.
- The ingress router to re-signal the LSP.

Configuring IS-IS 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.

Configuring IS-IS Traffic Engineering Attributes

You can configure the following IS-IS traffic engineering attributes:

- Configuring IS-IS to Use IGP Shortcuts on page 67
- Configuring IS-IS to Ignore the Metric of RSVP Label-Switched Paths on page 69
- Disabling IS-IS Support for Traffic Engineering on page 69
- Installing IPv4 Routes into the Multicast Routing Table on page 69
- Configuring IS-IS to Use Protocol Preference to Determine the Traffic Engineering Database Credibility Value on page 70

When configuring traffic engineering support, you can also configure IS-IS to use metric values greater than 63, as described in "Enabling Wide IS-IS Metrics for Traffic Engineering" on page 62.

Configuring IS-IS to Use IGP Shortcuts

IS-IS always performs shortest-path-first (SPF) calculations to determine next hops. For prefixes reachable through a particular next hop, IS-IS places that next hop for that prefix in the **inet.0** routing table. In addition, for routers running MPLS, IS-IS installs the

prefix for IPv4 routes in the **inet.3** routing table as well. The **inet.3** table, which is present on the ingress router, contains the host address of each MPLS label-switched path (LSP) egress router. BGP uses this routing table to resolve next-hop addresses.

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 label-switched paths as shortcuts when installing information in the **inet.3** or **inet6.3** routing table, include the following statements:

```
traffic-engineering {
  family inet {
    shortcuts;
  }
  shortcuts inet6 {
    shortcuts;
  }
}
```

For IPv4 traffic, include the **inet** statement. For IPv6 traffic, include the **inet6** statement.

To configure load balancing across multiple LSPs, include the **multipath** statement.

When traffic engineering shortcuts are used, RSVP first looks at the **metric2** value, which is derived from the IGP cost. After this, RSVP considers the LSP metric value. So, if a certain path changes for an LSP and the cost changes, not all LSPs are used to load-balance the network.

When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) combinations can be lost during route calculation. To override this default behavior for load balancing, include the **lsp-equal-cost** statement to retain the equal cost path information in the routing table.

```
traffic-engineering {
  multipath {
    lsp-equal-cost;
    }
```

```
}
```

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

Because the **inet.3** routing table is present only on ingress routers, you can configure LSP shortcuts only on these routers.

For more information about configuring LSPs and MPLS, see the *Junos OS MPLS Applications Configuration Guide*.

Configuring IS-IS to Ignore the Metric of RSVP Label-Switched Paths

You can configure IS-IS to ignore the metric of RSVP label-switched paths (LSPs) when LDP tunneling is enabled. If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops. Ignoring the metric of RSVP LSPs avoids mutual dependency between IS-IS and RSVP, eliminating the time period when the RSVP metric used for tunneling traffic is not up to date.

To configure IS-IS to ignore the metric of RSVP LSPs, include the **ignore-lsp-metrics** statement:

```
traffic-engineering {
  ignore-lsp-metrics;
}
```

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

For more information about configuring LSPs and MPLS, see the *Junos OS MPLS Applications Configuration Guide*.

Disabling IS-IS Support for Traffic Engineering

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:

```
traffic-engineering {
  disable;
}
```

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

Installing IPv4 Routes into the Multicast Routing Table

You can install unicast IPv4 routes into the multicast routing table (**inet.2**) for multicast reverse-path forwarding (RPF) checks.

To install routes into the multicast routing table for RPF checks, include the **multicast-rpf-routes** statement:

```
traffic-engineering {
  family inet {
    shortcuts {
      multicast-rpf-routes;
    }
  }
}
```

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

NOTE: Traffic engineering shortcuts must be enabled.

NOTE: IPv4 multicast topology must not be enabled.

NOTE: LSPs must not be advertised into IS-IS.

Configuring IS-IS to Use Protocol Preference to Determine the Traffic Engineering Database Credibility Value

By default, Junos OS prefers IS-IS routes in the traffic engineering database over other IGP routes even if the routes of another IGP are configured with a lower, that is, more preferred, preference value. 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. For more information, see Example: Enabling OSPF Traffic Engineering Support.

To configure IS-IS to use the configured protocol preference for IGP routes to determine the traffic engineering database credibility value, include the **credibility-protocol-preference** statement at the **[edit protocols isis traffic-engineering]** hierarchy level:

```
[edit protocols isis]
traffic-engineering {
  credibility-protocol-preference;
```

}

Enabling Authentication for IS-IS 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.

Configuring Quicker Advertisement of IS-IS Adjacency State Changes

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.

To disable the adjacency hold-down timer, include the **no-adjacency-holddown** statement:

```
no-adjacency-holddown;
```

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

Enabling Padding of IS-IS Hello Packets

You can configure padding on hello packets to accommodate asymmetrical maximum transfer units (MTUs) from different routing devices. This help prevents 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.

To configure padding for hello packets, include the **hello-padding** statement:

```
hello-padding (adaptive | loose | strict);
```

There are three 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 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. 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 may result in generation of

extra LSPs that are flooded throughout the network. Specify the **adaptive** option to configure enough padding to establish an adjacency to neighbors.

- 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 may 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.

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

Configuring CLNS for IS-IS

Connectionless Network Services (CLNS) is a Layer 3 protocol, similar to IP version 4 (IPv4). CLNS uses network service access points (NSAPs) to address end systems and intermediate systems.

You can use IS-IS as the IGP to carry ISO CLNS routes through a network.

NOTE: CLNS is supported on J Series Services Routers and MX Series routers only.

To enable IS-IS to exchange CLNS routes, include the **clns-routing** statement:

```
clns-routing;
```

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

You can configure a pure CLNS network by disabling IPv4 and IPv6 for IS-IS.

To disable IPv4, include the **no-ipv4-routing** statement:

```
no-ipv4-routing;
```

To disable IPv6, include the **no-ipv6-routing** statement:

```
no-ipv6-routing;
```

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

You can export BGP routes into Layer 2 IS-IS by configuring an export policy and applying the policy to IS-IS. You can export BGP routes from a specific VRF instance into IS-IS by configuring and applying an export policy at the **[edit routing-instance *instance-name* protocols isis]** hierarchy level. ES-IS routes from one routing instance cannot be exported into a Layer 1 IS-IS area of another routing instance.

To configure an export policy to export BGP routes into IS-IS, include the **policy-statement** statement:

```
policy-statement policy-name {
  from {
    protocol bgp;
    family iso;
  }
  then {
    accept;
  }
}
```

To apply an export policy, include the **export** statement at the **[edit protocols isis]** hierarchy level:

```
export policy-name;
```

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

For more information on policy configuration, see the *Junos OS Policy Framework Configuration Guide*.

You can also export routes from protocols other than BGP into IS-IS. ES-IS routes are exported to IS-IS by default. You can export ES-IS routes into IS-IS by configuring a routing policy.

Example: Configuring CLNS for IS-IS

Configure a routing policy to accept CLNS routes:

```
policy-options {
  policy-statement dist-bgp {
    from {
      protocol bgp;
      family iso;
    }
    then accept;
  }
  policy-statement dist-static {
    from {
      protocol static;
      family iso;
    }
    then accept;
  }
}
```

Configure CLNS for IS-IS:

```
protocols {
  isis {
    traceoptions {
      file isis size 5m world-readable;
      flag error;
```

```
        }
        export dist-static;
        no-ipv6-routing;
        no-ipv4-routing;
        clns-routing;
        interface fe-0/0/1.0;
        interface t1-0/2/1.0;
        interface fxp0.0 {
          disable;
        }
        interface lo0.0;
      }
    }
```

Configure a routing instance that supports CLNS routes:

```
routing-instances {
  aaaa {
    instance-type vrf;
    interface lo0.1;
    interface t1-3/0/0.0;
    interface fe-5/0/1.0;
    route-distinguisher 10.245.245.1:1;
    vrf-target target:11111:1;
    protocols {
      isis {
        export dist-bgp;
        no-ipv4-routing;
        no-ipv6-routing;
        clns-routing;
        interface all;
      }
    }
  }
}
```

Disabling IS-IS

To disable IS-IS on the routing device without removing the IS-IS configuration statements from the configuration, include the **disable** statement:

```
isis {
  disable;
}
```

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

To reenable IS-IS, remove the **disable** statement from the configuration:

```
[edit protocols]
user@host# delete isis disable
[edit protocols]
user@host# show
isis;
```

Disabling IPv4 Routing for IS-IS

You can disable IP version 4 (IPv4) routing for IS-IS. Disabling IPv4 routing results in the following:

- The routing device does not advertise the NLPID for IPv4 in the Junos OS link-state PDU fragment zero.
- The routing device does not advertise any IPv4 prefixes in Junos OS link-state PDUs.
- The routing device does not advertise the NLPID for IPv4 in Junos OS hello packets.
- The routing device does not advertise any IPv4 addresses in Junos OS hello packets.
- The routing device does not calculate any IPv4 routes.

To disable IPv4 routing on the routing device, include the **no-ipv4-routing** statement:

```
isis {
  no-ipv4-routing;
}
```

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

To reenable IS-IS, remove the **no-ipv4-routing** statement from the configuration:

```
[edit protocols]
user@host# delete isis no-ipv4-routing
```

NOTE: Note: Even when **no-ipv4-routing** is configured, an IS-IS traceoptions log can list rejected IPv4 addresses. When a configuration is committed, IS-IS schedules a scan of the routing table to determine whether any routes need to be exported into the IS-IS link state database. The implicit default export policy action is to reject everything. IPv4 addresses from the routing table are examined for export, rejected by the default policy, and the rejections are logged.

Disabling IPv6 Routing for IS-IS

You can disable IP version 6 (IPv6) routing for IS-IS. Disabling IPv6 routing results in the following:

- Routing device does not advertise the NLPID for IPv6 in Junos OS 0th link-state PDU fragment.
- Routing device does not advertise any IPv6 prefixes in Junos OS link-state PDUs.
- Routing device does not advertise the NLPID for IPv6 in Junos OS hello packets.
- Routing device does not advertise any IPv6 addresses in Junos OS hello packets.
- Routing device does not calculate any IPv6 routes.

To disable IPv6 routing on the routing device, include the **no-ipv6-routing** statement:

```
isis {
  no-ipv6-routing;
}
```

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

To re-enable IS-IS, remove the **disable** statement from the configuration:

```
[edit protocols]
user@host# delete isis no-ipv6-routing
```

Applying Policies to Routes Exported to IS-IS

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.

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 *routing policy* for that protocol. For information about defining routing policy, see the *Junos OS Policy Framework Configuration Guide*.

To apply routing policies that affect how the routing protocol process (rpd) exports routes into IS-IS, include the **export** statement:

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

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

NOTE: For IS-IS, you cannot apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol can easily lead to an inconsistent topology database.

Examples: Configuring IS-IS Routing Policy

Define a policy that allows only host routes from USC (**128.125.0.0/16**), and apply the policy to routes exported from the routing table into IS-IS:

```
policy-options {
  policy-statement usc-hosts-only {
    term first {
      from {
        route-filter 128.125.0.0/16 upto /31;
      }
      then reject;
    }
    then accept;
  }
```

```
}
protocols {
   isis {
      export usc-hosts-only;
   }
}
```

Define a policy that takes BGP routes from the **Edu** community and places them into IS-IS with a metric of 14. Apply the policy to routes exported from the routing table into IS-IS:

```
protocols {
   isis {
      export edu-to-isis;
   }
}
policy-options {
   community Edu members 666:5;
   policy-statement edu-to-isis {
      from {
         protocol bgp;
         community Edu;
      }
      to protocol isis;
      then metric 14;
   }
}
```

Define a policy that rejects all IS-IS Level 1 routes so that none are exported into IS-IS:

```
policy-options {
   policy-statement level1 {
      term first {
         from level 1;
         then reject;
      }
      then accept;
   }
}
protocols {
   isis {
      export level1;
      interface fxp0;
   }
}
```

Define a routing policy to export IS-IS Level 1 internal-only routes into Level 2:

```
[edit]
protocols {
   isis {
      export L1-L2;
   }
}
policy-statement L1-L2 {
   term one {
      from {
```

```
        level 1;
        external;
      }
      then reject;
    }
    term two {
      from level 1;
      to level 2;
      then accept;
    }
  }
```

Define a routing policy to export IS-IS Level 2 routes into Level 1:

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

Configuring 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 LSP paths.

The 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, the Junos OS creates for all destination routes that traverse a protected interface, a single alternate path to the primary next hop. 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, the 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 5 on page 80, Case 1 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 2, 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 5: Link Protection and Node-Link Protection Comparison for IS-IS Routes

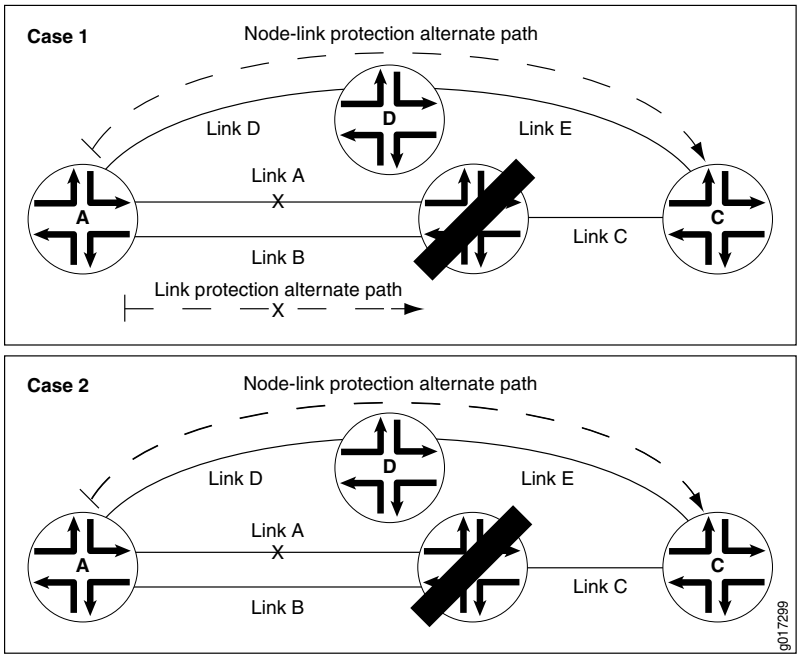

The Junos OS implementation of support for loop-free alternate paths for IS-IS routes is based on the following standards:

- Internet draft draft-ietf-rtgwg-ipfrr-spec-base-12.txt, *Basic Specification for IP Fast-Reroute: Loop-free Alternates*
- Internet draft draft-ietf-rtgwg-ipfrr-framework-06.txt, *IP Fast Reroute Framework*

This section discusses the following topics:

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, the 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 Configuration Guide*.

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 are protected.
- **show isis interface detail**—Displays the type of protection (link or node-link) applied to each protected interface.

For more detailed information about these commands, see the *Junos OS Routing Protocols and Policies Command Reference*.

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

In this example, all the logical interfaces on the router are enabled for IS-IS level 2, LDP, and RSVP. Node-link protection is enabled on all the interfaces, which means that if the primary next hop for any destination that traverses the interfaces becomes unavailable, the 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.

Configure the interfaces. Enable IS-IS and MPLS. In this example, the interfaces are also enabled for both IPv4 and IPv6 traffic.

```
[edit interfaces]
ge-2/0/0 {
    unit 0 {
        family inet {
            address 11.14.0.1/30;
        }
        family iso;
        family inet6;
        family mpls;
    }
}
ge-2/0/1 {
    unit 0 {
        family inet {
            address 11.14.1.1/30;
        }
        family iso;
        family inet6;
        family mpls;
    }
}
so-3/0/1 {
    unit 0 {
        family inet {
            address 11.16.1.1/30;
        }
        family iso;
        family inet6;
        family mpls;
    }
}
so-3/0/2 {
    unit 0 {
        family inet {
            address 11.16.0.1/30;
        }
        family iso;
```

```
      family inet6;
      family mpls;
    }
  }
  so-6/0/0 {
    unit 0 {
      family inet {
        address 11.12.0.1/30;
      }
      family iso;
      family inet6;
      family mpls;
    }
  }
```

Configure the IS-IS interfaces for Level 2 only, and configure MPLS to use both RSVP and LDP label-switched paths (LSPs). Enable IS-IS node-link protection, which also automatically extends backup coverage to all LDP LSPs.

```
[edit protocols]
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
isis {
  interface all {
    node-link-protection; # Enable node-link protection on all IS-IS interfaces.# Protection
      is automatically extended to all LDP LSPs.
    level 2 metric 10;
    level 1 disable;
  }
  interface fxp0.0 {
    disable;
  }
  interface lo0.0 {
    level 2 metric 0;
  }
}
ldp {
  deaggregate; # Enable forwarding equivalence class deaggregation, which results in
    faster global convergence.
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

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. After this policy is configured, export it to the forwarding table of the local router with the **export** statement at the **[edit routing-options forwarding-table]** hierarchy level.

```
[edit policy-options]
policy-statement ecmp {
  term 1 {
    then {
      load-balance per-packet;
    }
  }
}
```

```
[edit routing-options]
forwarding-table {
  export ecmp;
}
```

Related Documentation

- Example: Load Balancing BGP Traffic

Disabling Adjacency Down and Neighbor Down Notification in IS-IS and OSPF

Whenever IS-IS is deactivated, the IS-IS adjacencies are brought down. IS-IS signals to RSVP to bring down any RSVP neighbors associated with the IS-IS adjacencies, and this further causes the associated LSPs signaled by RSVP to go down as well.

A similar process occurs whenever OSPF is deactivated. The OSPF neighbors are brought down. OSPF signals to RSVP to bring down any of the RSVP neighbors associated with the OSPF neighbors, and this further causes the associated LSPs signaled by RSVP to go down as well.

If you need to migrate from IS-IS to OSPF or from OSPF to IS-IS, the IGP notification to RSVP for an adjacency or neighbor down event needs to be ignored. Using the **no-adjacency-down-notification** or **no-neighbor-down-notification** statements, you can disable IS-IS adjacency down notification or OSPF neighbor down notification, respectively, until the migration is complete. The network administrator is responsible for configuring the statements before the migration, and then removing them from the configuration afterward, so that IGP notification can function normally.

To disable adjacency down notification in IS-IS, include the **no-adjacency-down-notification** statement:

```
no-adjacency-down-notification;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols isis interface *interface-name*]**
- **[edit logical-systems *logical-system-name* protocols isis interface *interface-name*]**

To disable neighbor down notification in OSPF, include the **no-neighbor-down-notification** statement:

```
no-neighbor-down-notification;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols ospf area *area-id* interface *interface-name*]**
- **[edit logical-systems *logical-system-name* protocols ospf area *area-id* interface *interface-name*]**

Tracing IS-IS Protocol Traffic

You can trace various types of IS-IS protocol traffic to help debug IS-IS protocol issues. To trace IS-IS protocol traffic, include the **traceoptions** statement at the **[edit protocols isis]** hierarchy level:

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

You can specify the following IS-IS protocol-specific trace options using the **flag** statement:

- **csn**—Complete sequence number PDU (CSNP) packets
- **error**—Errored packets
- **graceful-restart**—Graceful restart operations
- **hello**—Hello packets
- **ldp-synchronization**—Synchronization between IS-IS and LDP
- **lsp**—Link-state PDU packets
- **lsp-generation**—Link-state PDU generation packets
- **nsr-synchronization**—NSR synchronization events
- **packets**—All IS-IS protocol packets
- **psn**—Partial sequence number PDU (PSNP) packets
- **spf**—Shortest-path-first (SPF) calculations

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

NOTE: Use the flag modifier **detail** with caution as this may cause the CPU to become very busy.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the IS-IS protocol using the **traceoptions flag** statement included at the **[edit protocols isis]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

NOTE: Use the trace flag **all** with caution as this may cause the CPU to become very busy.

Examples: Tracing IS-IS Protocol Traffic

A common configuration traces SPF calculations, LSP calculations, normal protocol operations, and errors in protocol operation:

```
[edit]
protocols {
   isis {
      traceoptions {
         file isis-log size 1m files 10;
         flag spf;
         flag lsp;
         flag error;
         flag normal;
      }
   }
}
```

Trace only unusual or abnormal operations to the file **routing-log**, and trace detailed information about all IS-IS packets to the file **isis-log**:

```
[edit]
routing-options {
   traceoptions {
      file routing-log;
   }
}
protocols {
   isis {
      traceoptions {
```

```
        file isis-log size 10k files 5;
        flag csn detail;
        flag hello detail;
        flag lsp detail;
        flag psn detail;
      }
    }
  }
```

Perform detailed tracing of mesh-group flooding:

```
[edit]
protocols {
  isis {
    traceoptions {
      file isis-log;
      flag lsp detail;
    }
  }
}
```

IS-IS LSP packets that contain errors are discarded by default. To log these errors, specify the **error** tracing operation:

```
[edit]
protocols {
  isis {
    traceoptions {
      file isis-log;
      flag error;
    }
  }
}
```

Related Documentation

- traceoptions on page 166
- For more information about tracing and global tracing options, see Example: Tracing Global Routing Protocol Operations.

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

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

Requirements

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

Overview

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

Figure 6 on page 89 shows the sample network.

Figure 6: IS-IS on Logical Systems

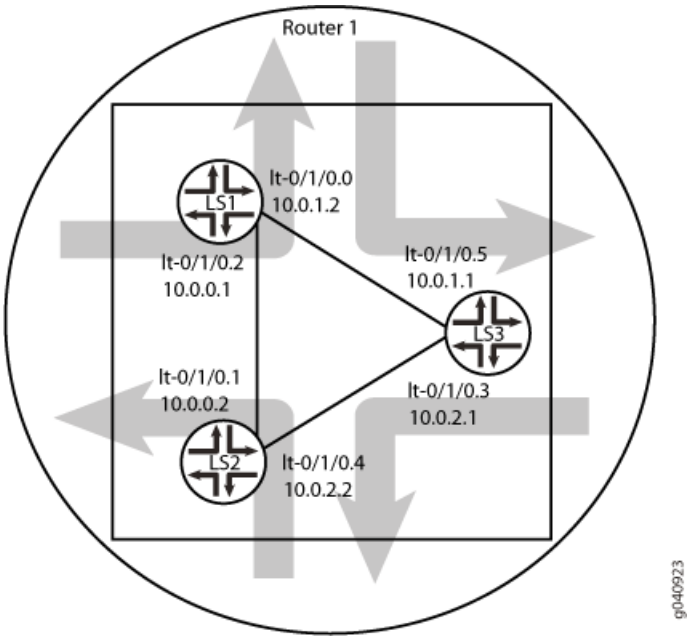

Configuration

CLI Quick Configuration

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

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

```
set logical-systems LS1 interfaces lo0 unit 1 family iso address 49.0001.1720.1600.1001.00
set logical-systems LS1 protocols isis interface lt-0/1/0.0
set logical-systems LS1 protocols isis interface lt-0/1/0.2
set logical-systems LS1 protocols isis interface lo0.1 passive
set logical-systems LS2 interfaces lt-0/1/0 unit 1 description LS2->LS1
set logical-systems LS2 interfaces lt-0/1/0 unit 1 encapsulation ethernet
set logical-systems LS2 interfaces lt-0/1/0 unit 1 peer-unit 2
set logical-systems LS2 interfaces lt-0/1/0 unit 1 family inet address 10.0.0.2/30
set logical-systems LS2 interfaces lt-0/1/0 unit 1 family iso
set logical-systems LS2 interfaces lt-0/1/0 unit 4 description LS2->LS3
set logical-systems LS2 interfaces lt-0/1/0 unit 4 encapsulation ethernet
set logical-systems LS2 interfaces lt-0/1/0 unit 4 peer-unit 3
set logical-systems LS2 interfaces lt-0/1/0 unit 4 family inet address 10.0.2.2/30
set logical-systems LS2 interfaces lt-0/1/0 unit 4 family iso
set logical-systems LS2 interfaces lo0 unit 2 family iso address
  49.0001.1720.1600.2002.00
set logical-systems LS2 protocols isis interface lt-0/1/0.1
set logical-systems LS2 protocols isis interface lt-0/1/0.4
set logical-systems LS2 protocols isis interface lo0.2 passive
set logical-systems LS3 interfaces lt-0/1/0 unit 3 description LS3->LS2
set logical-systems LS3 interfaces lt-0/1/0 unit 3 encapsulation ethernet
set logical-systems LS3 interfaces lt-0/1/0 unit 3 peer-unit 4
set logical-systems LS3 interfaces lt-0/1/0 unit 3 family inet address 10.0.2.1/30
set logical-systems LS3 interfaces lt-0/1/0 unit 3 family iso
set logical-systems LS3 interfaces lt-0/1/0 unit 5 description LS3->LS1
set logical-systems LS3 interfaces lt-0/1/0 unit 5 encapsulation ethernet
set logical-systems LS3 interfaces lt-0/1/0 unit 5 peer-unit 0
set logical-systems LS3 interfaces lt-0/1/0 unit 5 family inet address 10.0.1.1/30
set logical-systems LS3 interfaces lt-0/1/0 unit 5 family iso
set logical-systems LS3 interfaces lo0 unit 3 family iso address 49.0001.1234.1600.2231.00
set logical-systems LS3 protocols isis interface lt-0/1/0.5
set logical-systems LS3 protocols isis interface lt-0/1/0.3
set logical-systems LS3 protocols isis interface lo0.3 passive
```

Step-by-Step Procedure

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

To configure IS-IS on logical systems:

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

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

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

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

```
user@host# set interfaces lt-0/1/0 unit 0 peer-unit 5
user@host# set interfaces lt-0/1/0 unit 0 family inet address 10.0.1.2/30
user@host# set interfaces lt-0/1/0 unit 0 family iso
```

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

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

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

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

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

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

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

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

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

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

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

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

8. Configure IS-IS on all the interfaces.

```
[edit logical-systems LS1 protocols isis]
```

```
user@host# set interface lt-0/1/0.0
user@host# set interface lt-0/1/0.2

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

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

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

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

Results Confirm your configuration by issuing the **show logical-systems** command.

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

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

        }
    }
}
LS3 {
    interfaces {
        lt-0/1/0 {
            unit 3 {
                description LS3->LS2;
                encapsulation ethernet;
                peer-unit 4;
                family inet {
                    address 10.0.2.1/30;
                }
                family iso;
            }
            unit 5 {
                description LS3->LS1;
                encapsulation ethernet;
                peer-unit 0;
                family inet {
                    address 10.0.1.1/30;
                }
                family iso;
            }
        }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying That the Logical Systems Are Up on page 94
- Verifying Connectivity Between the Logical Systems on page 94

Verifying That the Logical Systems Are Up

Purpose Make sure that the interfaces are properly configured.

Action

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

Verifying Connectivity Between the Logical Systems

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

Action

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

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

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

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


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

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

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

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

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

10.0.0.0/30        *[IS-IS/15] 3w0d 01:37:10, metric 20
                      to 10.0.2.2 via lt-0/1/0.3
                    > to 10.0.1.2 via lt-0/1/0.5
10.0.1.0/30        *[Direct/0] 3w0d 01:38:10
                    > via lt-0/1/0.5
10.0.1.1/32        *[Local/0] 3w0d 01:38:11
                      Local via lt-0/1/0.5
10.0.2.0/30        *[Direct/0] 3w0d 01:38:11
                    > via lt-0/1/0.3
10.0.2.1/32        *[Local/0] 3w0d 01:38:11
                      Local via lt-0/1/0.3

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

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

From LS1, Ping LS3

```
user@host> set cli logical-system LS1

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

From LS3, Ping LS1

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

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

From LS1, Ping LS2

```
user@host> set cli logical-system LS1

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

From LS2, Ping LS1

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

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

From LS2, Ping LS3

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

user@host:LS2> ping 10.0.1.1
PING 10.0.1.1 (10.0.1.1): 56 data bytes
64 bytes from 10.0.1.1: icmp_seq=0 ttl=64 time=1.253 ms
64 bytes from 10.0.1.1: icmp_seq=1 ttl=64 time=1.194 ms
64 bytes from 10.0.1.1: icmp_seq=2 ttl=64 time=1.212 ms
64 bytes from 10.0.1.1: icmp_seq=3 ttl=64 time=1.221 ms
64 bytes from 10.0.1.1: icmp_seq=4 ttl=64 time=1.195 ms
^C
```

```
--- 10.0.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.194/1.215/1.253/0.022 ms
```

From LS3, Ping LS2

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

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

Related Documentation

- Example: Creating an Interface on a Logical System
- Example: Connecting Logical Systems Within the Same Router Using Logical Tunnel Interfaces

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

Before you begin:

- Connect the logical systems by using logical tunnel (**lt**) interfaces. See Example: Connecting Logical Systems Within the Same Router Using Logical Tunnel Interfaces.
- Enable IS-IS on the interfaces. See "Example: Configuring IS-IS on Logical Systems Within the Same Router" on page 88.

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 98 shows the sample network.

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

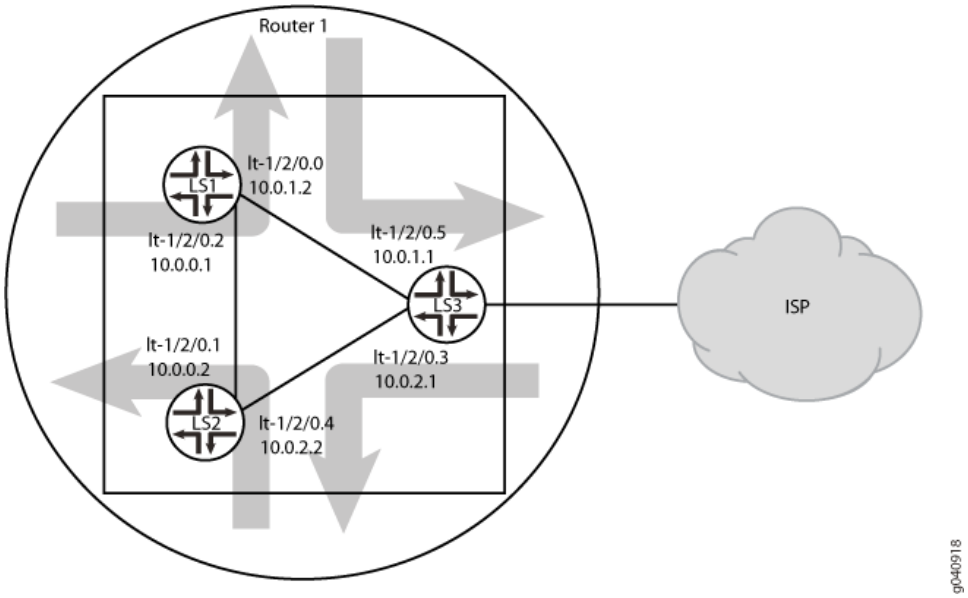

Configuration

CLI Quick Configuration

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

```
set logical-systems LS3 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
set logical-systems LS3 protocols isis export isis-default
```

Step-by-Step Procedure

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

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

1. Configure the default route on Logical System LS3.

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

2. Configure the default route policy on Logical System LS3.

   ```
   [edit logical-systems LS3 policy-options]
   user@host# set policy-statement isis-default from protocol static
   ```

```
user@host# set policy-statement isis-default from route-filter 0.0.0.0/0 exact
user@host# set policy-statement isis-default then accept
```

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

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

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

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

Results Confirm your configuration by issuing the **show logical-systems LS3** command.

```
user@host# show logical-systems LS3
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@host> 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@host> 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@host> show route logical-system LS1
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[IS-IS/160] 00:02:01, metric 10
                    > to 10.0.1.1 via lt-1/2/0.0
10.0.0.0/30        *[Direct/0] 1w0d 10:16:34
                    > via lt-1/2/0.2
10.0.0.1/32        *[Local/0] 1w0d 10:16:34
                      Local via lt-1/2/0.2
10.0.1.0/30        *[Direct/0] 1w0d 10:16:34
                    > via lt-1/2/0.0
10.0.1.2/32        *[Local/0] 1w0d 10:16:34
                      Local via lt-1/2/0.0
10.0.2.0/30        *[IS-IS/15] 1w0d 10:15:55, metric 20
                      to 10.0.1.1 via lt-1/2/0.0
                    > to 10.0.0.2 via lt-1/2/0.2

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

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

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

Related Documentation

- Example: Creating an Interface on a Logical System

PART 3

Administration

CHAPTER 4

IS-IS Reference

IS-IS Standards

IS-IS is defined in the following documents:

- ISO 8473, *Protocol for providing the connectionless-mode network services*
- ISO 9542, *End System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for the Provision of the Connectionless-mode Network Service*
- ISO 10589, *Intermediate System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for the Provision of the Connectionless-mode Network Service*
- RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*
- RFC 2973, *IS-IS Mesh Groups*
- RFC 3787, *Recommendations for Interoperable IP Networks Using Intermediate System to Intermediate System (IS-IS)*
- RFC 5120, *M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)*
- RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*
- RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*
- RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*
- RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*
- RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*
- RFC 5304, *IS-IS Cryptographic Authentication*
- RFC 5305, *IS-IS Extensions for Traffic Engineering*
- RFC 5306, *Restart Signaling for IS-IS*
- RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*

- RFC 5308, *Routing IPv6 with IS-IS*
- RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
- Internet draft draft-ietf-bfd-base-09.txt, *Bidirectional Forwarding Detection* (except for the transmission of echo packets)
- Internet draft draft-ietf-isis-wg-255adj-02.txt, *Maintaining more than 255 circuits in IS-IS*

To access Internet RFCs and drafts, go to the Internet Engineering Task Force (IETF) Web site at **http://www.ietf.org**.

CHAPTER 5

Summary of IS-IS Configuration Statements

authentication-key

Syntax authentication-key *key*;

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

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

Description Authentication key (password). Neighboring routing devices use the password to verify the authenticity of packets sent from this interface. For the key to work, you also must include the **authentication-type** statement.

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.

Default If you do not include this statement and the **authentication-type** statement, IS-IS authentication is disabled.

Options *key*—Authentication password. The password can be up to 1024 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

CAUTION: A simple password for authentication is truncated if it exceeds 254 characters.

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

Related Documentation

- Configuring IS-IS Authentication on page 19

authentication-key-chain

Syntax authentication-key-chain *key-chain-name*;

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

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Apply and enable an authentication keychain to the routing device.

Options ***key-chain***—Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

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

Related Documentation

- Example: Configuring Route Authentication for BGP
- Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols
- Example: Configuring BFD Authentication for Static Routes
- Overview of Hitless Authentication Key Rollover for IS-IS on page 10
- Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 21

authentication-type

Syntax authentication-type *authentication*;

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

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

Description Enable authentication and specify the authentication scheme for IS-IS. If you enable authentication, you must specify a password by including the **authentication-key** statement.

Default If you do not include this statement and the **authentication-key** statement, IS-IS authentication is disabled.

Options ***authentication***—Authentication scheme:

- **md5**—Use HMAC authentication in combination with MD5. HMAC-MD5 authentication is defined in RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*.
- **simple**—Use a simple password for authentication. The password is included in the transmitted packet, making this method of authentication relatively insecure. We recommend that you *not* use this authentication method.

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

Related Documentation

- authentication-key on page 108
- no-authentication-check on page 149
- Configuring IS-IS Authentication on page 19

bfd-liveness-detection

Syntax

```
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;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}
```

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.
detection-time threshold and **transmit-interval threshold** options added in Junos OS Release 8.2.
Support for logical systems introduced in Junos OS Release 8.3.
no-adaptation statement introduced in Junos OS Release 9.0.
authentication algorithm, **authentication key-chain**, and **authentication loose-check** options introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure bidirectional failure detection timers and authentication.

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 may not be configured at both ends of the BFD session.

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

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.

Required Privilege Level

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

Related Documentation
- Configuring BFD for IS-IS on page 27
- Configuring BFD Authentication for IS-IS on page 35

checksum

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

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

Related Documentation
- Enabling Packet Checksums on IS-IS Interfaces on page 38

context-identifier

Syntax context-identifier *identifier*

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis],
[edit protocols isis]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure IS-IS context identifier information.

Options *identifier*—IPv4 address that defines a protection pair. The context identifier is manually configured on both primary and protector PEs.

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

Related Documentation
- show isis context-identifier

clns-routing

Syntax clns-routing;

Hierarchy Level [edit protocols isis],
[edit routing-instances *routing-instance-name* protocols isis]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable IS-IS to exchange CLNS routes.

NOTE: CLNS is supported on J Series Services Routers and MX Series routers only.

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

Related Documentation

- Configuring CLNS for IS-IS on page 72

csnp-interval

Syntax csnp-interval (*seconds* | disable);

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 Configure the interval between complete sequence number (CSN) packets on a LAN interface.

Options **disable**—Do not send CSN packets on this interface.

seconds—Number of seconds between the sending of CSN packets.

Range: 1 through 65,535 seconds

Default: 10 seconds

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

Related Documentation

- Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces on page 38

disable (IS-IS)

Syntax disable;

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

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 Disable IS-IS on the routing device, on an interface, or on a level. At the **[edit protocols isis traffic-engineering]** hierarchy level, disable IS-IS support for traffic engineering.

Enabling IS-IS on an interface (by including the **interface** statement at the **[edit protocols isis]** or the **[edit routing-instances *routing-instance-name* protocols isis]** hierarchy level), disabling it (by including the **disable** statement), and not actually having IS-IS run on an interface (by including the **passive** statement) are mutually exclusive states.

Default IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which an International Organization for Standardization (ISO) protocol family is enabled.

IS-IS support for traffic engineering is enabled.

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

Related Documentation
- IS-IS Overview on page 3
- Disabling IS-IS Support for Traffic Engineering on page 69
- Disabling IS-IS on page 74

disable (LDP Synchronization)

Syntax disable;

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 7.5.

Description Disable LDP for IS-IS.

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

Related Documentation
- Configuring Synchronization Between LDP and IS-IS on page 38

export

Syntax export [*policy-names*];

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

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

Description Apply one or more policies to routes being exported from the routing table into IS-IS.

Options ***policy-names***—Name of one or more policies.

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

Related Documentation
- Applying Policies to Routes Exported to IS-IS on page 76
- *Junos OS Policy Framework Configuration Guide*

external-preference

Syntax external-preference *preference*;

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

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

Description Configure the preference of external routes.

Options *preference*—Preference value.

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

Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)

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

Related Documentation

- preference on page 157
- Configuring Preference Values for IS-IS Routes on page 62

family

Syntax

```
family inet {
  shortcuts {
    multicast-rpf-routes;
  }
}
family inet6 {
  shortcuts;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols isis level],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis level],
[edit protocols isis level],
[edit routing-instances *routing-instance-name* protocols isis level]

Release Information

Statement introduced in Junos OS Release 9.3.
Support for IPv6 for IGP shortcuts introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure the address family for traffic engineering IS-IS interior gateway protocol (IGP) shortcuts.

Options

inet—IPv4 address family

inet6—IPv6 address family

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring IS-IS Traffic Engineering Attributes on page 67

graceful-restart

Syntax

```
graceful-restart {
  disable;
  helper-disable;
  restart-duration seconds;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis],
[edit protocols isis]

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

Description Configure graceful restart for IS-IS.

Options

disable—Disable graceful restart.

helper-disable—Disable graceful restart helper capability. Helper mode is enabled by default.

restart-duration *seconds*—Configure the time period for the restart to last, in seconds.

Range: 30 through 300 seconds

Default: 30 seconds

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

Related Documentation

- *Junos OS High Availability Configuration Guide*

hello-authentication-key

Syntax hello-authentication-key *password*;

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

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

Description Configure an authentication key (password) for hello packets. Neighboring routing devices use the password to verify the authenticity of packets sent from an interface. For the key to work, you also must include the **hello-authentication-type** statement.

Default By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.

Options *password*—Authentication password. The password can be up to 255 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

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

Related Documentation

- authentication-key on page 108
- authentication-type on page 110
- hello-authentication-type on page 123
- Configuring Authentication for IS-IS Hello Packets on page 58

hello-authentication-key-chain

Syntax hello-authentication-key-chain *key-chain-name*;

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

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Apply an authentication keychain to the IS-IS interface.

Options ***key-chain-name***—Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

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

Related Documentation

- Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 21
- Overview of Hitless Authentication Key Rollover for IS-IS on page 10

hello-authentication-type

Syntax hello-authentication-type (md5 | simple);

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

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

Description Enable authentication on an interface for hello packets. If you enable authentication on hello packets, you must specify a password by including the **hello-authentication-key** statement.

Default By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.

Options
md5—Specifies Message Digest 5 as the packet verification type.

simple—Specifies simple authentication as the packet verification type.

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

Related Documentation
- authentication-key on page 108
- authentication-type on page 110
- hello-authentication-key on page 121
- Configuring Authentication for IS-IS Hello Packets on page 58

hello-interval

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.

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

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)

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

Related Documentation
- hold-time on page 127
- Configuring the Transmission Frequency for IS-IS Hello Packets on page 59

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.

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

Options
adaptive—Configure padding until the state of the neighbor adjacency is up.

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

strict—Configure padding for all adjacency states.

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

Related Documentation
- Enabling Padding of IS-IS Hello Packets on page 71

hold-time (IS-IS)

Syntax hold-time *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.

Description 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.

Options
seconds—Hold-time value, in seconds.

Range: 3 through 65,535 seconds, or 1 to send out hello packets every 333 milliseconds

Default: 9 seconds (for designated intermediate system [DIS] routers), 27 seconds (for non-DIS routers; three times the default hello interval)

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

Related Documentation
- hello-interval on page 124
- Configuring the Delay Before IS-IS Neighbors Mark the Routing Device as Down on page 59

hold-time (LDP Synchronization)

Syntax hold-time *seconds*;

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

Release Information Statement introduced in Junos OS Release 7.5.

Description Configure the time period to advertise the maximum cost metric for a link that is not fully operational.

Options ***seconds***—Hold-time value, in seconds.

Range: 1 through 65,535 seconds

Default: Infinity

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

Related Documentation

- Configuring Synchronization Between LDP and IS-IS on page 38

ignore-attached-bit

Syntax ignore-attached-bit;

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

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

Description 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 LSPs. If the attached bit is ignored, no default route, which points to the routing device which has set the attached bit, is installed.

Default The **ignore-attached-bit** statement is disabled by default.

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

Related Documentation
- Configuring IS-IS on page 16

ignore-lsp-metrics

Syntax ignore-lsp-metrics;

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

Release Information Statement introduced in Junos OS Release 8.0.

Description Ignore the metrics for RSVP label-switched paths in IS-IS traffic engineering shortcut calculations or when you configure LDP over RSVP label-switched paths.

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

Related Documentation
- shortcuts on page 162
- Configuring IS-IS to Use IGP Shortcuts on page 67
- *Junos OS MPLS Applications Configuration Guide*

interface

Syntax

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

Hierarchy Level

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

Release Information

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

Description

Configure interface-specific IS-IS properties. To configure more than one interface, include the **interface** statement multiple times.

Enabling IS-IS on an interface (by including the **interface** statement at the **[edit protocols isis]** or the **[edit routing-instances *routing-instance-name* protocols isis]** hierarchy level), disabling it (by including the **disable** statement), and not actually having IS-IS run on an interface (by including the **passive** statement) are mutually exclusive states.

Options

all—Have Junos OS create IS-IS interfaces automatically.

interface-name—Name of an interface. Specify the full interface name, including the physical and logical address components.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring of Interface-Specific IS-IS Properties on page 26

ipv4-multicast

Syntax ipv4-multicast;

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

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

Description Configure alternate IPv4 multicast topologies.

Default Multicast topologies are disabled.

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

Related Documentation
- Configuring IS-IS Multicast Topology on page 40

ipv4-multicast-metric

Syntax ipv4-multicast-metric *metric*;

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.

Description Specify the multicast topology metric value for the level.

Options *metric*—Metric value.
Range: 0 through 16,777,215

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

Related Documentation
- Configuring IS-IS Multicast Topology on page 40

ipv6-multicast

Syntax ipv6-multicast;

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

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

Description Configure alternate IPv6 multicast topologies.

Default Multicast topologies are disabled.

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

Related Documentation
- Configuring IS-IS Multicast Topology on page 40

ipv6-multicast-metric

Syntax ipv6-multicast-metric *metric*;

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.

Description Specify the IPv6 alternate multicast topology metric value for the level.

Options *metric*—Metric value.
Range: 0 through 16,777,215

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

Related Documentation
- Configuring IS-IS Multicast Topology on page 40

ipv6-unicast

Syntax ipv6-unicast;

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

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

Description Configure alternate IPv6 unicast topologies.

Default IPv6 unicast topologies are disabled.

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

Related Documentation
- Configuring IS-IS IPv6 Unicast Topologies on page 55

ipv6-unicast-metric

Syntax ipv6-unicast-metric *metric*;

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.

Description Specify the IPv6 unicast topology metric value for the level.

Options *metric*—Metric value.
Range: 0 through 16,777,215

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

Related Documentation
- Configuring IS-IS IPv6 Unicast Topologies on page 55

isis

Syntax isis { ... }

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

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

Description Enable IS-IS routing on the routing device or for a routing instance.

The **isis** statement is the one statement you must include in the configuration to run IS-IS on the routing device or in a routing instance.

Default IS-IS is disabled on the routing device.

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

Related Documentation

- Minimum IS-IS Configuration on page 18

label-switched-path

Syntax label-switched-path *name* level *level-number* metric *metric*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Advertise LSPs into IS-IS as point-to-point links. The LSP is advertised in the appropriate IS-IS levels as a point-to-point link and contains a local address and a remote address.

Options ***name***—Identifies the LSP.

level-number—IS-IS level number.

Values: 1 or 2

metric—Metric value.

Range: 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)

Default: 0 (for **lo0**), 10 (for all other interfaces)

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

Related Documentation

- Advertising Label-Switched Paths into IS-IS on page 63

LDP-synchronization

Syntax

```
ldp-synchronization {
  disable;
  hold-time seconds;
}
```

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 7.5.

Description

Enable synchronization by advertising the maximum cost metric until LDP is operational on the link.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Example: Configuring Synchronization Between LDP and IGPs

level (Global IS-IS)

Syntax

```
level level-number {
    authentication-key key;
    authentication-key-chain key-chain-name;
    authentication-type type;
    external-preference preference;
    no-csnp-authentication;
    no-hello-authentication;
    no-psnp-authentication;
    preference preference;
    wide-metrics-only;
}
```

Hierarchy Level

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

Release Information

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

Description

Configure the global-level properties.

Options

level-number—IS-IS level number.

Values: 1 or 2

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Preference Values for IS-IS Routes on page 62

level (IS-IS Interfaces)

Syntax

```
level level-number {
  level (IS-IS Interfaces);
  hello-authentication-key key;
  hello-authentication-key-chain key-chain-name;
  hello-authentication-type authentication;
  hello-interval seconds;
  hold-time seconds;
  ipv4-multicast-metric number;
  ipv6-unicast-metric number;
  metric metric;
  passive;
  priority number;
  te-metric metric;
}
```

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.

Description

Configure the IS-IS level. You can configure one instance of Level 1 routing and one instance of Level 2 routing on each interface, and you can configure the two levels differently.

Options

level-number—IS-IS level number.

Values: 1 or 2

Default: The routing device operates as both a Level 1 and Level 2 router.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Levels on IS-IS Interfaces on page 56

link-protection

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.

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
- node-link-protection on page 154
- Configuring Link Protection for IS-IS on page 80

loose-authentication-check

Syntax loose-authentication-check;

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

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

Description Allow the use of MD5 authentication without requiring network-wide deployment.

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

Related Documentation
- Enabling Authentication for IS-IS Without Network-Wide Deployment on page 71

lsp-equal-cost

Syntax lsp-equal-cost;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis traffic-engineering multipath], [edit protocols isis traffic-engineering multipath]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure label-switched paths (LSPs) to be retained as equal cost paths for load balancing when a better path metric is found during the IS-IS internal routing table calculation.

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.

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

Related Documentation

- Configuring IS-IS Traffic Engineering Attributes on page 67
- multipath on page 147 (IS-IS)
- traffic-engineering on page 169

lsp-interval

Syntax lsp-interval *milliseconds*;

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 Configure the link-state PDU interval time.

Options ***milliseconds***—Number of milliseconds between the sending of link-state PDUs. Specifying a value of 0 blocks all link-state PDU transmission.

Range: 0 through 1000 milliseconds

Default: 100 milliseconds

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

Related Documentation

- Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces on page 39

lsp-lifetime

Syntax lsp-lifetime *seconds*;

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

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

Description 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.

Options ***seconds***—link-state PDU lifetime, in seconds.

Range: 350 through 65,535 seconds

Default: 1200 seconds

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

Related Documentation

- Configuring the Link-State PDU Lifetime for IS-IS on page 63

max-areas

Syntax max-areas *number*;

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

Release Information
Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Modify the maximum number of IS-IS areas advertised.

Options
number—Maximum number of areas to include in the IS-IS hello (IIH) PDUs and link-state PDUs.

Range: 3 through 36

Default: 3

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

Related Documentation
- Limiting the Number of Advertised IS-IS Areas on page 62

mesh-group

Syntax mesh-group (blocked | *value*);

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 Configure an interface to be part of a mesh group, which is a set of fully connected nodes.

Options **blocked**—Configure the interface so that it does not flood link-state PDU packets.

value—Number that identifies the mesh group.

Range: 1 through 4,294,967,295 (2^{32} – 1; 32 bits are allocated to identify a mesh group)

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

Related Documentation

- Configuring Mesh Groups of IS-IS Interfaces on page 39

metric

Syntax metric *metric*;

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.

Description Specify the metric value for the level.

Options
metric—Metric value.
Range: 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)
Default: 10 (for all interfaces except **lo0**), 0 (for the **lo0** interface)

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

Related Documentation
- te-metric on page 164
- wide-metrics-only on page 170
- Configuring the Metric Value for IS-IS Routes on page 60

multicast-rpf-routes

Syntax multicast-rpf-routes;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis traffic-engineering family inet shortcuts],
[edit logical-systems *logical-system-name* routing-instances traffic-engineering family inet shortcuts],
[edit protocols isis traffic-engineering family inet shortcuts],
[edit routing-instances *routing-instance-name* protocols isis traffic-engineering family inet shortcuts]

Release Information Statement introduced in Junos OS Release 9.3.

Description Install IPv4 routes into the multicast routing table for RPF checks.

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

Related Documentation
- Installing IPv4 Routes into the Multicast Routing Table on page 69

multipath

Syntax

```
multipath {
  lsp-equal-cost;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols isistraffic-engineering], [edit protocols isistraffic-engineering]

Release Information Statement introduced in Junos OS Release 9.6.

Description Enable load balancing for multiple label-switched paths (LSPs).

Options

lsp-equal-cost—Configure LSPs to be retained as equal cost paths for load balancing when a better route metric is added to the routing table.

When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) path combinations can be lost during route calculation. To override this default IS-IS behavior, include the **lsp-equal-cost** statement for load balancing, so that the equal cost path information is retained in the routing table.

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

Related Documentation

- Configuring IS-IS Traffic Engineering Attributes on page 67
- lsp-equal-cost on page 140
- multipath (BGP)
- multipath (protocol independent)
- traffic-engineering on page 169

no-adjacency-down-notification

Syntax no-adjacency-down-notification;

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

Release Information Statement introduced in Junos OS Release 8.0.

Description Disable adjacency down notification for IS-IS to allow for migration from IS-IS to OSPF without disruption of the RSVP neighbors and associated RSVP-signaled LSPs.

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

Related Documentation
- Disabling Adjacency Down and Neighbor Down Notification in IS-IS and OSPF on page 85

no-adjacency-holddown

Syntax no-adjacency-holddown;

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

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Disable the hold-down timer for IS-IS adjacencies.

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

Related Documentation
- Configuring Quicker Advertisement of IS-IS Adjacency State Changes on page 71

no-authentication-check

Syntax no-authentication-check;

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

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

Description Generate authenticated packets and check the authentication on received packets, but do not reject packets that cannot be authenticated.

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

Related Documentation
- csnp-interval on page 115
- hello-authentication-type on page 123
- Configuring IS-IS Authentication on page 19

no-csnp-authentication

Syntax no-csnp-authentication;

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

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

Description Suppress authentication check on complete sequence number PDU (CSNP) packets.

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

Related Documentation
- csnp-interval on page 115
- Configuring IS-IS Authentication on page 19

no-eligible-backup

Syntax no-eligible-backup;

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

Description Exclude the specified interface as a backup interface for IS-IS interfaces on which link protection or node-link protection is enabled.

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

Related Documentation

- link-protection on page 139
- node-link-protection on page 154
- Excluding an IS-IS Interface as a Backup for Protected Interfaces on page 81

no-hello-authentication

Syntax no-hello-authentication;

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

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

Description Suppress authentication check on complete sequence number hello packets.

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

Related Documentation

- hello-authentication-type on page 123
- Configuring IS-IS Authentication on page 19

no-ipv4-multicast

Syntax no-ipv4-multicast;

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 Exclude an interface from IPv4 multicast topologies.

Default Multicast topologies are disabled.

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

Related Documentation
- Configuring IS-IS Multicast Topology on page 40

no-ipv4-routing

Syntax no-ipv4-routing;

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

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

Description Disable IP version 4 (IPv4) routing.

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

Related Documentation
- Disabling IPv4 Routing for IS-IS on page 75

no-ipv6-multicast

Syntax no-ipv6-multicast;

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.

Description Exclude an interface from the IPv6 multicast topologies.

Default Multicast topologies are disabled.

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

Related Documentation
- Configuring IS-IS Multicast Topology on page 40

no-ipv6-routing

Syntax no-ipv6-routing;

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

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

Description Disable IP version 6 (IPv6) routing.

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

Related Documentation
- Disabling IPv6 Routing for IS-IS on page 75

no-ipv6-unicast

Syntax no-ipv6-unicast;

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.

Description Exclude an interface from the IPv6 unicast topologies.

Default IPv6 unicast topologies are disabled.

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

Related Documentation
- Configuring IS-IS IPv6 Unicast Topologies on page 55

no-psnp-authentication

Syntax no-psnp-authentication;

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

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

Description Suppress authentication check on partial sequence number PDU (PSNP) packets.

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

Related Documentation
- Configuring IS-IS Authentication on page 19

no-unicast-topology

Syntax no-unicast-topology;

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 Exclude an interface from the IPv4 unicast topologies.

Default IPv4 unicast topologies are disabled.

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

Related Documentation
- Configuring IS-IS Multicast Topology on page 40

node-link-protection

Syntax node-link-protection;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis interface *interface-name*],
[edit logical-routers *logical-router-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 9.5 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Enable node-link protection on the specified IS-IS interface. Junos OS creates an alternate loop-free path to the primary next hop for all destination routes that traverse a protected interface. This alternate path avoids the primary next-hop routing device altogether and establishes a path through a different routing device.

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

Related Documentation
- link-protection on page 139
- Configuring Node-Link Protection for IS-IS on page 81

overload

Syntax

```
overload {
  advertise-high-metrics;
  allow-route-leaking;
  timeout seconds;
}
```

Hierarchy Level

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

Release Information

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

Description

Configure the local routing device so that it appears to be overloaded. You might want to do this when you want the routing device to participate in IS-IS routing, but do not want it to be used for transit traffic. Note that traffic to immediately attached interfaces continues to transit the routing device. You can also advertise maximum link metrics in network layer reachability information (NLRI) instead of setting the overload bit.

NOTE: If the time elapsed after the IS-IS instance is enabled is less than the specified timeout, overload mode is set.

Options

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

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

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

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

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

Default: 0 seconds

Range: 60 through 1800 seconds

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

Related Documentation
- Configuring IS-IS to Make Routing Devices Appear Overloaded on page 64

passive

Syntax
passive;

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

Description
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.

This statement effectively prevents IS-IS from running on the interface. To enable IS-IS on an interface, include the **interface** statement at the **[edit protocols isis]** or the **[edit routing-instances *routing-instance-name* protocols isis]** hierarchy level. To disable it, include the **disable** statement at those hierarchy levels. The three states—enabling, disabling, or not running IS-IS on an interface—are mutually exclusive.

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

Related Documentation
- disable on page 117
- Advertising Interface Addresses Without Running IS-IS on page 58

point-to-point

Syntax point-to-point;

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 Configure an IS-IS interface to behave like a point-to-point connection.

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

Related Documentation
- Configuring Point-to-Point Interfaces for IS-IS on page 56

preference

Syntax preference *preference*;

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

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

Description Configure the preference of internal routes.

Options ***preference***—Preference value.

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

Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)

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

Related Documentation
- external-preference on page 118
- Configuring Preference Values for IS-IS Routes on page 62

prefix-export-limit

Syntax prefix-export-limit *number*;

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

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

Description Configure a limit to the number of prefixes exported into IS-IS.

Options *number*—Prefix limit.

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

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

Related Documentation

- Limiting the Number of Prefixes Exported to IS-IS on page 63

priority

Syntax priority *number*;

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.

Description The interface's priority for becoming the designated router. The interface with the highest priority value becomes that level's designated router.

The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.

Options *number*—Priority value.

Range: 0 through 127

Default: 64

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

Related Documentation
- Configuring the Designated Router Priority for IS-IS on page 61

reference-bandwidth

Syntax reference-bandwidth *reference-bandwidth*;

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

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

Description Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula:

```
cost = reference-bandwidth/bandwidth
```

Options ***reference-bandwidth***—Reference bandwidth value in bits per second.

Default: None

Range: 9600 through 1,000,000,000,000 bps

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

Related Documentation

- Configuring the Reference Bandwidth Used in IS-IS Metric Calculations on page 61

rib-group

Syntax

```
rib-group {
    inet group-name;
    inet6 group-name;
}
```

Hierarchy Level

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

Release Information

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

Description

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.

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.

Options

group-name—Name of the routing table group.

inet—Install IPv4 IS-IS routes.

inet6—Install IPv6 IS-IS routes.

Required Privilege Level

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

Related Documentation

- Example: Exporting Specific Routes from One Routing Table Into Another Routing Table
- Example: Importing Direct and Static Routes Into a Routing Instance
- Understanding Multiprotocol BGP

shortcuts

Syntax

```
shortcuts {
    multicast-rpf-routes;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols isis traffic-engineering family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis traffic-engineering family (inet | inet6)],
[edit protocols isis traffic-engineering family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols isis traffic-engineering family (inet | inet6)]

Release Information

Statement introduced before Junos OS Release 7.4.
The **family** statement and support for IPv6 routes for IS-IS traffic engineering shortcuts introduced in Junos OS Release 9.3.

Description

Configure IS-IS to use MPLS label-switched paths (LSPs) as next hops if possible when installing routing information into the **inet.3** or **inet6.3** routing table.

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- Configuring IS-IS Traffic Engineering Attributes on page 67

spf-options

Syntax

```
spf-options {
    delay milliseconds;
    holddown milliseconds;
    rapid-runs number;
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description

Configure options for running the shortest-path-first (SPF) algorithm. You can configure a delay for when to run the SPF algorithm after a network topology change is detected, the maximum number of times the SPF algorithm can run in succession, and a holddown interval after SPF algorithm runs the maximum number of times.

Options

delay *milliseconds*—Time interval between the detection of a topology change and when the SPF algorithm runs.

Range: 50 through 1000 milliseconds

Default: 200 milliseconds

holddown *milliseconds*—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.

Range: 2000 through 10,000 milliseconds

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

Required Privilege Level

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

Related Documentation

- Configuring SPF Options for IS-IS on page 65

te-metric

Syntax te-metric *metric*;

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.

Description Metric value used by traffic engineering for information injected into the traffic engineering database. The value of the traffic engineering metric does not affect normal IS-IS forwarding.

Options *metric*—Metric value.

Range: 1 through 16,777,215

Default: Value of the IGP metric

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

Related Documentation

- metric on page 145
- wide-metrics-only on page 170
- Configuring the Metric Value for IS-IS Routes on page 60

topologies

Syntax

```
topologies {
  ipv4-multicast;
  ipv6-multicast;
  ipv6-unicast;
}
```

Hierarchy Level

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

Release Information

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

Description

Configure alternate IS-IS topologies.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring IS-IS Multicast Topology on page 40

traceoptions

Syntax

```
traceoptions {
    file name <size size> <files number> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

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

Release Information

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

Description

Configure IS-IS protocol-level tracing options. To specify more than one tracing operation, include multiple **flag** statements.

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

Default

The default IS-IS protocol-level tracing options are those inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *name*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks (“ ”). All files are placed in the directory **/var/log**. We recommend that you place IS-IS tracing output in the file **isis-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

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
- **ldp-synchronization**—Synchronization between IS-IS and LDP
- **lsp**—Link-state PDU packets
- **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: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

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

Related Documentation

- Tracing IS-IS Protocol Traffic on page 86

traffic-engineering

Syntax

```
traffic-engineering {
  credibility-protocol-preference {
    disable;
    family inet {
      shortcuts {
        multicast-rpf-routes;
      }
    }
    family inet6 {
      shortcuts;
    }
    multipath {
      lsp-equal-cost;
    }
  }
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Support for the **family** statement introduced in Junos OS Release 9.3.
Support for the **credibility-protocol-preference** statement introduced in Junos OS Release 9.4.
Support for the **multipath** statement introduced in Junos OS Release 9.6.
Support for the **lsp-equal-cost** statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure traffic engineering properties for IS-IS.

Default

IS-IS traffic engineering support is enabled.

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.

multipath—Enable load balancing for multiple label-switched paths (LSPs).

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.

The remaining statements are explained separately.

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

Related Documentation
- Configuring IS-IS Traffic Engineering Attributes on page 67
- traffic-engineering

wide-metrics-only

Syntax
wide-metrics-only;

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

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

Description
Configure IS-IS to generate metric values greater than 63 on a per IS-IS level basis.

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

Related Documentation
- te-metric on page 164
- Enabling Wide IS-IS Metrics for Traffic Engineering on page 62

PART 4

Troubleshooting

CHAPTER 6

Routing Protocol Process Memory FAQ

Routing Protocol Process Memory FAQ Overview

The Juniper Networks Junos operating system (Junos OS) is based on the FreeBSD Unix operating system. The open source software is modified and hardened to operate in the device's specialized environment. For example, some executables have been deleted while other utilities have been de-emphasized. Additionally, certain software processes have been added to enhance the routing functionality. The result of this transformation is the kernel, the heart of the Junos OS software.

The kernel is responsible for generating multiple processes that perform the actual functions of the device. Each process operates in its own protected memory space, providing isolation between the processes and resiliency in the event of a process failure. This is important in a core routing platform because a single process failure does not cause the entire device to cease functioning.

Some of the common software processes include the routing protocol process (rpd) that controls the device's protocols, the device control process (dcd) that controls the device's interfaces, the management process (mgd) that controls user access to the device, the chassis process (chassisd) that controls the device's properties itself, and the Packet Forwarding Engine process (pfed) that controls the communication between the device's Packet Forwarding Engine and the Routing Engine. Besides the above processes, there are other specialized processes that support additional functionality, such as the Simple Network Management Protocol (SNMP), Virtual Router Redundancy Protocol (VRRP), and Class of Service (CoS).

The routing protocol process is a software process within the Routing Engine software that controls the routing protocols that run on the device. Its functionality includes all protocol messages, routing table updates, and implementation of routing policies.

The routing protocol process starts all configured routing protocols and handles all routing messages. It maintains one or more routing tables, which consolidate the routing information learned from all routing protocols. From this routing information, the routing protocol process determines the active routes to network destinations and installs these routes into the Routing Engine's forwarding table. Finally, it implements the routing policy, which allows you to control the routing information that is transferred between the routing

protocols and the routing table. Using the routing policy, you can filter and limit the transfer of information as well as set properties associated with specific routes.

Related Documentation

- Routing Protocol Process Memory FAQs on page 174

Routing Protocol Process Memory FAQs

The following sections present the most frequently asked questions and answers related to the routing protocol process memory utilization, operation, interpretation of related command outputs, and troubleshooting the software process.

Routing Protocol Process Memory Utilization FAQs

This section presents frequently asked questions and answers related to the memory usage of the routing protocol process.

Why does the routing protocol process use excessive memory?

The routing protocol process uses hundreds of megabytes of RAM in the Routing Engine to store information needed for the operation of routing and related protocols, such as BGP, OSPF, ISIS, RSVP, LDP, and MPLS. Such huge consumption of memory is common for the process, as the information it stores includes routes, next hops, interfaces, routing policies, labels, and label-switched paths (LSPs). Because access to the RAM memory is much faster than access to the hard disk, most of the routing protocol process information is stored in the RAM memory instead of using the hard disk space. This ensures that the performance of the routing protocol process is maximized.

How can I check the amount of memory the routing protocol process is using?

You can check the routing protocol process memory usage by entering the **show system processes** and the **show task memory** Junos OS command-line interface (CLI) operational mode commands.

The **show system processes** command displays information about software processes that are running on the device. You can check the routing protocol process memory usage by using the **show system processes** command with the **extensive** option.

The **show task memory** command displays a report generated by the routing protocol process on the memory utilization for routing protocol tasks on the Routing Engine. Although the report generated by the routing protocol process is on its own memory usage, it does not display all the memory used by the process. The value reported by the routing protocol process does not account for the memory used for the **TEXT** and **STACK** segments, or the memory used by the process's internal memory manager. The **show task memory** command also does not include the memory which has been deactivated by the routing protocol process, although some or all of that deactivated memory has not actually been freed by the kernel.

For more information about checking the routing protocol process memory usage, see Check Routing Protocol Process (rpd) Memory Usage in the *Junos OS Baseline Network Operations Guide*.

For more information about the show system processes command and the show task memory command, see the *Junos OS System Basics and Services Command Reference*.

I just deleted many routes from the routing protocol process. Why is the routing protocol process still using so much memory?

The **show system processes extensive** command displays a **RES** value measured in kilobytes. This value represents the amount of process memory resident in the physical memory. This is also known as RSS or Resident Set Size. Any amount of memory deactivated by the process might still be considered part of the **RES** value. Generally, the kernel defers the actual freeing of deactivated memory until there is a memory shortage. This can lead to large discrepancies between the values reported by the routing protocol process and the kernel, even after the routing protocol process has deactivated a large amount of memory.

Interpreting Routing Protocol Process-Related Command Outputs FAQs

This section presents frequently asked questions and answers about the routing protocol process-related Junos OS CLI command outputs that are used to display the memory usage of the routing protocol process.

How do I interpret memory numbers displayed in the show system processes extensive command output?

The **show system processes extensive** command displays exhaustive system process information about software processes that are running on the device. This command is equivalent to the UNIX **top** command. However, the UNIX **top** command shows real-time memory usage, with the memory values constantly changing, while the **show system processes extensive** command provides a snapshot of memory usage in a given moment.

To check overall CPU and memory usage, enter the **show system processes extensive** command. Refer to Table 7 on page 177 for information about the **show system processes extensive** command output fields.

```
user@host> show system processes extensive
last pid:   544;  load averages:  0.00,  0.00,  0.00    18:30:33
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3968K Inact, 19M Wired, 184K Cache, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free
  PID USERNAME PRI NICE SIZE    RES STATE    TIME   WCPU    CPU COMMAND
  544 root      30   0   604K   768K RUN      0:00  0.00%  0.00% top
    3 root      28   0     0K    12K psleep   0:00  0.00%  0.00% vmdaemon
    4 root      28   0     0K    12K update   0:03  0.00%  0.00% update
  528 aviva     18   0   660K   948K pause    0:00  0.00%  0.00% tcsh
  204 root      18   0   300K   544K pause    0:00  0.00%  0.00% csh
  131 root      18   0   332K   532K pause    0:00  0.00%  0.00% cron
  186 root      18   0   196K    68K pause    0:00  0.00%  0.00% watchdog
   27 root      10   0   512M 16288K mfsidl   0:00  0.00%  0.00% mount_mfs
    1 root      10   0   620K   344K wait     0:00  0.00%  0.00% init
  304 root       3   0   884K   900K ttyin    0:00  0.00%  0.00% bash
  200 root       3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  203 root       3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  202 root       3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  201 root       3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  194 root       2   0  2248K  1640K select   0:11  0.00%  0.00% rpd
  205 root       2   0   964K   800K select   0:12  0.00%  0.00% tnp.chassisd
  189 root       2 -12   352K   740K select   0:03  0.00%  0.00% xntpd
  114 root       2   0   296K   612K select   0:00  0.00%  0.00% amd
```

```
  188 root      2    0   780K   600K select   0:00  0.00%  0.00% dcd
  527 root      2    0   176K   580K select   0:00  0.00%  0.00% rlogind
  195 root      2    0   212K   552K select   0:00  0.00%  0.00% inetd
  187 root      2    0   192K   532K select   0:00  0.00%  0.00% tnetd
   83 root      2    0   188K   520K select   0:00  0.00%  0.00% syslogd
  538 root      2    0  1324K   516K select   0:00  0.00%  0.00% mgd
   99 daemon    2    0   176K   492K select   0:00  0.00%  0.00% portmap
  163 root      2    0   572K   420K select   0:00  0.00%  0.00% nsrexecd
  192 root      2    0   560K   400K select   0:10  0.00%  0.00% snmpd
  191 root      2    0  1284K   376K select   0:00  0.00%  0.00% mgd
  537 aviva     2    0   636K   364K select   0:00  0.00%  0.00% cli
  193 root      2    0   312K   204K select   0:07  0.00%  0.00% mib2d
    5 root      2    0     0K    12K pfesel   0:00  0.00%  0.00% if_pfe
    2 root    -18    0     0K    12K psleep   0:00  0.00%  0.00% pagedaemon
    0 root    -18    0     0K     0K sched    0:00  0.00%  0.00% swapper
```

Table 7 on page 177 describes the output fields that represent the memory values for the **show system processes extensive** command. Output fields are listed in the approximate order in which they appear.

Table 7: show system processes extensive Output Fields

Field Name	Field Description
Mem	Information about physical and virtual memory allocation.
Active	Memory allocated and actively used by the process.
Inact	Memory allocated but not recently used, or memory deactivated by the processes. Inactive memory remains mapped in the address space of one or more processes and, therefore, counts toward the RSS value of those processes.
Wired	Memory that is not eligible to be swapped, usually used for in-kernel memory structure, memory physically locked by a process, or both.
Cache	Freed memory that is no longer associated with any process but still has valid contents that correspond to some file system blocks. Cache pages can be reclaimed as is when the corresponding file system blocks are accessed again. However, when the system is under memory pressure, the contents of Cache pages could be erased by the kernel and the pages reused to service any memory allocation requests.
Buf	Size of the virtual memory buffer used to hold data recently called from the disk.
Free	Free memory that is neither associated with any process nor contains any valid contents.
Swap	Information about swap memory. • Total—Total space on the swap device. • Used—Memory swapped to disk. • Free—Unused space available on the swap device.

The rest of the command output displays information about the memory usage of each process. The **SIZE** field indicates the size of the virtual address space, and the **RES** field indicates the amount of the process in physical memory, which is also known as RSS or Resident Set Size. For more information, see the show system processes command in the *Junos OS System Basics and Services Command Reference*.

What is the difference between Active and Inact memory that is displayed by the show system processes extensive command?

When the system is under memory pressure, the pageout process can free up memory from the **Inact** and, if necessary, **Active** pools after first preserving the contents of those pages on the swap device or backing file systems if necessary. When the pageout process runs, it scans memory to see which pages are good candidates to be unmapped and freed up. Thus, the distinction between **Active** and **Inact** memory is only used by the pageout process to determine which pool of pages to free first at the time of a memory shortage.

The pageout process first scans the **Inact** list and checks whether the pages on this list have been accessed since the time they have been listed here. The pages that have been accessed are moved from the **Inact** list to the **Active** list. On the other hand, pages that have not been accessed become prime candidates to be freed by the pageout process. If the pageout process cannot produce enough free pages from the **Inact** list, pages from the **Active** list are freed up.

Because the pageout process runs only when the system is under memory pressure, the pages on the **Inact** list remain untouched – even if they have not been accessed recently – when the amount of **Free** memory is adequate.

How do I interpret memory numbers displayed in the show task memory command output?

The **show task memory** command provides a comprehensive picture of the memory utilization for routing protocol tasks on the Routing Engine. The routing protocol process is the main task that uses Routing Engine memory.

To check routing process memory usage, enter the **show task memory** command.

```
user@host> show task memory
Memory                 Size (kB)  %Available  When
  Currently In Use:        29417          3%  now
  Maximum Ever Used:       33882          4%  00/02/11 22:07:03
  Available:              756281        100%  now
```

Table 8 on page 178 describes the output fields for the **show task memory** command. Output fields are listed in the approximate order in which they appear.

Table 8: show task memory Output Fields

Field Name	Field Description
Memory Currently In Use	Memory currently in use. Dynamically allocated memory plus the **DATA** segment memory in kilobytes.
Memory Maximum Ever Used	Maximum memory ever used.
Memory Available	Memory currently available.

The **show task memory** command does not display all the memory used by the routing protocol process. This value does not account for the memory used for the **TEXT** and

STACK segments, or the memory used by the routing protocol process's internal memory manager. The **show task memory** command also does not include the memory which has been deactivated by the routing protocol process, although some or all of that deactivated memory has not actually been freed by the kernel.

Why is the Memory Currently In Use value less than the RES value?

The **show task memory** command displays a **Memory Currently In Use** value measured in kilobytes. This value is the dynamically allocated memory plus the **DATA** segment memory. The **show system processes extensive** command displays a **RES** value measured in kilobytes. This value represents the amount of process memory resident in the physical memory. This is also known as RSS or Resident Set Size.

The **Memory Currently In Use** value does not account for all of the memory that the routing protocol process uses. This value does not include the memory used for the **TEXT** and the **STACK** segments, and a small percentage of memory used by the routing protocol process's internal memory manager. The **show task memory** command also does not include the memory which has been deactivated by the routing protocol process, although some or all of that deactivated memory has not actually been freed by the kernel.

Any amount of memory deactivated by the routing protocol process might still be considered part of the **RES** value. Generally, the kernel defers the actual freeing of deactivated memory until there is a memory shortage This can lead to large discrepancies between the **Memory Currently In Use** value and the **RES** value.

Routing Protocol Process Memory Swapping FAQs

This section presents frequently asked questions and answers related to the memory swapping of the routing protocol process from the Routing Engine memory to the hard disk memory.

Why does the system start swapping when I try to perform a core dump using the request system core-dumps command?

The **request system core-dumps** command displays a list of system core files created when the device has failed. This command can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification date, path, and filename. You can use the **core-filename** option and the **core-file-info**, **brief**, and **detail** options to display more information about the specified core dump files.

You can use the **request system core-dumps** command to perform a non-fatal core dump without aborting the routing protocol process. To do this, the routing protocol process is forked, generating a second copy, and then aborted. This process can double the memory consumed by the two copies of the routing protocol process, pushing the system into swap.

Why does the show system processes extensive command show that memory is swapped to disk even though there is plenty of free memory?

Memory can remain swapped out indefinitely if it is not accessed again. Therefore, the **show system processes extensive** command shows that memory is swapped to disk even though there is plenty of free memory. Such a situation is not unusual.

Troubleshooting the Routing Protocol Process FAQs

This section presents frequently asked questions and answers related to a shortage of memory and memory leakage by the routing protocol process.

What does the RPD_OS_MEMHIGH message mean?

The **RPD_OS_MEMHIGH** message is written into the system message file if the routing protocol process is running out of memory. This message alerts you that the routing protocol process is using the indicated amount and percentage of Routing Engine memory, which is considered excessive. This message is generated either because the routing protocol process is leaking memory or the use of system resources is excessive, perhaps because routing filters are not configured properly or the configured network topology is very complex.

When the memory utilization for the routing protocol process is using all available Routing Engine DRAM memory or reaches the maximum memory limit, a message of the following form is written every minute in the syslog message file:

```
RPD_OS_MEMHIGH: Using 188830 KB of memory, 100 percent of available
```

This message includes the amount (in kilobytes), the percentage, or both of the available memory in use.

This message should not appear under normal conditions, as any further memory allocations usually require a portion of existing memory to be written to swap. As a recommended solution, increase the amount of RAM in the Routing Engine. For more information, see http://kb.juniper.net/InfoCenter/index?page=content&id=KB14186 .

What can I do when there is a memory shortage even after a swap?

We do not recommend that the system operate in this state, notwithstanding the existence of swap. The protocols that run in the routing protocol process usually have a real-time requirement that cannot reliably withstand the latency of being swapped to hard disk. If the memory shortage has not resulted from a memory leak, then either a reduction in the memory usage or an upgrade to a higher memory-capacity Routing Engine is required.

What is the task_timer?

The source of a routing protocol process memory leak can usually be identified by dumping the timers for each task. You can use the **show task *task-name*** command to display routing protocol tasks on the Routing Engine. Tasks can be baseline tasks performed regardless of the device's configuration, and other tasks that depend on the device configuration.

For more information, see the show task command in the *Junos OS System Basics and Services Command Reference*.

Related Documentation

- Routing Protocol Process Memory FAQ Overview on page 173

PART 5

Index

Index

Symbols

A

B

C

D

E

K

L

M

N

O

P

R

S

T

V

W