

Chapter 12

IS-IS Configuration Guidelines

To configure Intermediate System-to-Intermediate System (IS-IS), you include the following statements in the configuration:

```
protocols {
  isis {
    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-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;
    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-delay milliseconds;
  }
}
```

```

topologies {
  ipv4-multicast;
  ipv6-multicast;
  ipv6-unicast;
}
traffic-engineering {
  disable;
  shortcuts;
}
traceoptions {
  file name <replace> <size size> <files number> <no-stamp>
    <(world-readable | no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}
interface interface-name {
  disable;
  bfd-liveness-detection {
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier number;
  }
  checksum;
  csnp-interval (seconds | disable);
  lsp-interval milliseconds;
  mesh-group (value | blocked);
  no-ipv4-multicast;
  no-ipv6-unicast;
  passive;
  point-to-point;
  level level-number {
    disable;
    hello-authentication-key key;
    hello-authentication-type authentication;
    hello-interval seconds;
    hold-time seconds;
    ipv4-multicast-metric number;
    ipv6-multicast-metric number;
    ipv6-unicast-metric number;
    metric metric;
    passive;
    priority number;
    te-metric metric;
  }
}
}
}

```

For a list of hierarchy levels at which you can configure 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 Standards Organization (ISO) address is configured.

This chapter discusses the following topics that provide information about configuring IS-IS:

- Minimum IS-IS Configuration on page 228
- Configuring IS-IS Authentication on page 228
- Configuring Interface-Specific Properties on page 230
- Enabling Checksum on page 231
- Configuring the CSNP Interval on page 231
- Configuring Mesh Groups on page 231
- Modifying the Interface Metric on page 232
- Enabling Wide Metrics for Traffic Engineering on page 232
- Configuring Route Preferences on page 233
- Configuring a Prefix Export Limit on page 233
- Configuring IS-IS Levels on an Interface on page 233
- Modifying the LSP Interval on page 239
- Modifying the LSP Lifetime on page 239
- Advertising Label-Switched Paths into IS-IS on page 239
- Configuring the Router to Appear Overloaded on page 240
- Configuring the SPF Delay on page 240
- Configuring Graceful Restart on page 241
- IS-IS and Multipoint Configurations on page 241
- Configuring Point-to-Point Interfaces on page 241
- Configuring IS-IS Traffic Engineering Attributes on page 242
- Configuring the BFD Protocol on page 243
- Configuring Loose Authentication Check on page 244
- Disabling IS-IS on page 244
- Disabling IPv4 Routing on page 244
- Disabling IPv6 Routing on page 245
- Configuring IS-IS Routing Policy on page 246
- Configuring IS-IS Multicast Topologies on page 248

Configuring IS-IS IPv6 Unicast Topologies on page 251

Installing Default Route to Nearest Level 1/Level 2 Router on page 251

Tracing IS-IS Protocol Traffic on page 252

Minimum IS-IS Configuration

For IS-IS to run on the router, you must enable IS-IS on the router, configure a network entity title (NET) on one of the router'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 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;
      }
    }
  }
  type-fpc/pic/port {
    unit logical-unit-number {
      family iso;
    }
  }
}
protocols {
  isis {
    isis all;
  }
}

```

Configuring IS-IS Authentication

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

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 router 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”).

HMAC-MD5 authentication—Uses an iterated cryptographic hash function. The receiving router 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 Hello Packets” on page 236.

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 configure this statement, see the statement summary section for this statement.

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

```
authentication-key key;
```

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

```
[edit protocols isis]
no-authentication-check;
```

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

```
[edit protocols isis]
no-hello-authentication;
```

To suppress authentication of PSNP packets, include the no-psnp-authentication statement:

```
[edit protocols isis]
no-psnp-authentication;
```

To suppress authentication of CSNP packets, include the no-csnp-authentication statement:

```
[edit protocols isis]
no-csnp-authentication;
```

For a list of hierarchy levels at which you can configure 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 interface hierarchy level and configuring no-authentication at the isis hierarchy level has no effect.

Configuring Interface-Specific Properties

You can configure interface-specific IS-IS properties by including the interface statement. These properties are explained later in this chapter.

```
[edit protocols isis]
interface interface-name {
  disable;
  bfd-liveness-detection {
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier number;
  }
  checksum;
  csnp-interval (seconds | disable);
  lsp-interval milliseconds;
  mesh-group (value | blocked);
  no-ipv4-multicast;
  no-ipv6-multicast;
  no-ipv6-unicast;
  passive;
  level level-number {
    disable;
    hello-authentication-type authentication;
    hello-authentication-key key;
    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 configure 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 information about configuring interfaces, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

Enabling Checksum

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

```
[edit protocols isis interface interface-name]  
checksum;
```

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

Configuring the CSNP Interval

By default, IS-IS sends complete sequence number (CSN) packets periodically. If the router is the designated router on a LAN, IS-IS sends CSN packets every 10 seconds. If the router 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:

```
[edit protocols isis interface interface-name]  
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:

```
[edit protocols isis interface interface-name]  
csnp-interval disable;
```

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

Configuring Mesh Groups

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

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

```
[edit protocols isis interface interface-name]  
mesh-group value;
```

To prevent an interface in the mesh group from flooding LSPs, configure blocking on that interface:

```
[edit protocols isis interface interface-name]  
mesh-group blocked;
```

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

Modifying the Interface Metric

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:

$$\text{cost} = \text{reference-bandwidth} / \text{bandwidth}$$

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).

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

```
[edit protocols isis]
reference-bandwidth reference-bandwidth;
```

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

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 interface metrics, see “Modifying the IS-IS Metric” on page 237.

Enabling Wide 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 $2^{24} - 1$ (16,777,215).

By default, the JUNOS software supports the sending and receiving of wide metrics. The JUNOS software 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:

```
[edit protocols isis level]
wide-metrics-only;
```

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

Configuring Route Preferences

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” on page 6.

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:

```
[edit protocols isis level level-number]
external-preference preference;
preference preference;
```

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

The preference value can range from 0 through 255.

Configuring a Prefix Export Limit

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:

```
[edit protocols isis level level-number]
prefix-export-limit number;
```

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

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

Configuring IS-IS Levels on an Interface

You can administratively divide a single AS into smaller groups called areas. You configure each router interface to be in an area. Any interface can be in any area. The area address applies to the entire router; 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 router can be a Level 1 router, a Level 2 router, or both. You specify the router level on a per-interface basis, and a router becomes adjacent with other routers 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:

```
[edit protocols isis interface interface-name]
level level-number {
  disable;
  hello-authentication-key key;
  hello-authentication-type authentication;
  hello-interval seconds;
  hold-time seconds;
  ipv4-multicast-metric number;
  ipv6-multicast-metric number;
  ipv6-unicast-metric number;
  metric metric;
  passive;
  priority number;
  te-metric metric;
}
```

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

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

Disabling IS-IS on a Level on page 234

Advertising Interface Addresses without Running IS-IS on page 235

Configuring Authentication for Hello Packets on page 236

Modifying the Hello Interval on page 237

Modifying the Hold-Time Value on page 237

Modifying the IS-IS Metric on page 237

Modifying the Traffic Engineering Metric on page 238

Configuring the Priority for Becoming the Designated Router on page 238

Configuring the Router to Advertise without Running IS-IS on page 238

Disabling IS-IS on a Level

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:

```
[edit protocols isis interface interface-name level level-number]
disable;
```

For a list of hierarchy levels at which you can configure 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 on a Level

On SONET/SDH interface so-0/0/0, enable IS-IS for Level 1 only. With this configuration, tracing messages periodically will indicate that IS-IS is creating Level 2 LSPs. However, because IS-IS for Level 2 is disabled, these LSPs 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;
```

For a list of hierarchy levels at which you can configure 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.



NOTE: If neither passive mode nor family ISO are configured on the IS-IS interface, then the router treats the interface as not being operational and no direct IPv4/IPv6 routes are exported into IS-IS.

Configuring Authentication for Hello Packets

You can configure authentication for all IS-IS hello packets for an interface and, to achieve a more fine-grained authentication, you can configure authentication for a given IS-IS level on that interface. If you configure a point-to-point link and if you enable both levels, the hello packets are sent with the password configured for Level 1.



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

To enable hello authentication for an interface, include the `hello-authentication-type` statement:

```
[edit protocols isis interface interface-name]  
hello-authentication-type authentication;
```

To configure the password, include the `hello-authentication-key` statement:

```
[edit protocols isis interface interface-name]  
hello-authentication-key key;
```

To enable hello authentication at an IS-IS level on an interface, include the `hello-authentication-type` statement:

```
[edit protocols isis interface interface-name level level-number]  
hello-authentication-type authentication;
```

To configure a password at an IS-IS level on an interface, include the `hello-authentication-key` statement:

```
[edit protocols isis interface interface-name level level-number]  
hello-authentication-key key;
```

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

Modifying the Hello Interval

Routers 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 intersystem (DIS) router sends hello packets every 3 seconds, and a non-DIS router sends hello packets every 9 seconds.

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

```
[edit protocols isis interface interface-name level level-number]
hello-interval seconds;
```

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

You can send out hello packets in sub-second 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 configure this statement, see the statement summary section for this statement.

Modifying the Hold-Time Value

The hold time specifies how long a neighbor should consider this router to be operative without receiving another hello packet. If the neighbor does not receive a hello packet from this router within the hold time, it marks the router as being unavailable. The default hold-time value is three times the default hello interval: 9 seconds for a DIS router and 27 seconds for a non-DIS router.

To modify the hold-time value on the local router, include the hold-time statement:

```
[edit protocols isis interface interface-name level level-number]
hold-time seconds;
```

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

Modifying the IS-IS Metric

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 $2^{24}-1$ (16,777,215) 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:

```
[edit protocols isis interface interface-name level level-number]
metric metric;
```

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

For more information about IS-IS interface metrics, see “Modifying the Interface Metric” on page 232.

Modifying the Traffic Engineering Metric

When traffic engineering is enabled on the router, 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 (TED). Its value does not affect normal IS-IS forwarding.

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

```
[edit protocols isis interface interface-name level level-number]
te-metric metric;
```

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

Configuring the Priority for Becoming the Designated Router

A router 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 router is responsible for sending network link-state advertisements, which describe all the routers attached to the network. These advertisements are flooded throughout a single area.

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

A router's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127; routers with a higher value are more likely to become the designated router. By default, routers have a priority value of 64.

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

```
[edit protocols isis interface interface-name level level-number]
priority number;
```

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

Configuring the Router to Advertise without Running IS-IS

The router can advertise the direct interface addresses on an interface or on a sub-level of the interface without actually running IS-IS on that interface or at that level. This occurs in passive mode.

To enable an interface as passive, include the `passive` statement:

```
[edit protocols isis interface interface-name level level-number]
passive;
```

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

Modifying the LSP Interval

By default, the router sends one LSP packet out an interface every 100 milliseconds. To modify this interval, include the `lsp-interval` statement:

```
[edit protocols isis interface interface-name]  
lsp-interval milliseconds;
```

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

To disable the transmission of all LSP packets, set the interval to 0.

Modifying the LSP Lifetime

By default, link-state PDUs (LSPs) 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 LSPs never expire.

To modify the LSP lifetime, include the `lsp-lifetime` statement:

```
[edit protocols isis]  
lsp-lifetime seconds;
```

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

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

The LSP refresh interval is derived from the LSP 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:

```
[edit protocols isis]  
label-switched-path name level level metric metric;
```

For a list of hierarchy levels at which you can configure 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.

For more information about advertising label-switched paths, see the *JUNOS Software MPLS Applications Configuration Guide*.

Configuring the Router to Appear Overloaded

You can configure the local router so that it appears to be overloaded. You might want to do this when you want the router 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 router.) To mark the router as overloaded, include the overload statement:

```
[edit protocols isis]
overload {
  advertise-high-metrics;
  <timeout seconds>;
}
```

To advertise maximum link metrics in NLRIs instead of setting the overload bit, include the advertise-high-metrics option when specifying the overload statement:

```
[edit protocols isis]
advertise-high-metrics;
```

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

```
[edit protocols isis]
overload timeout <seconds>;
```

The time can range from 60 through 1800 seconds.

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

Configuring the SPF Delay

You can configure the shortest-path-first (SPF) algorithm delay. The SPF algorithm delay is the amount of time between the detection of a topology change and when the SPF algorithm actually runs to achieve convergence. The shorter the delay, the shorter the convergence time.

To configure the SPF delay, include the `spf-delay` statement:

```
[edit protocols isis]
spf-delay milliseconds;
```

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

The time can range from 50 through 1000 milliseconds.

Configuring Graceful Restart

Graceful restart allows a router 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 router 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:

```
[edit protocols isis]
graceful-restart {
  helper-disable;
  restart-duration seconds;
}
```

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

To disable graceful restart for IS-IS, specify the `disable` statement. 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.

IS-IS and Multipoint Configurations

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 Point-to-Point Interfaces

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 will continue to treat the interface as a LAN interface. Only two IS-IS routers 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:

```
[edit protocols isis interface interface-name]
point-to-point;
```

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

Configuring IS-IS Traffic Engineering Attributes

You can configure two IS-IS traffic engineering attributes:

Configuring IS-IS to Use IGP Shortcuts on page 242

Disabling IS-IS Support for Traffic Engineering on page 243

When configuring traffic engineering support, you can also configure IS-IS to use metric values greater than 63, as described in “Enabling Wide Metrics for Traffic Engineering” on page 232.

Configuring IS-IS to Use IGP Shortcuts

IS-IS always performs 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 also installs the prefix in the inet.3 routing table. 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.



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

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 label-switched path as a next hop. The net result is that for BGP egress routers for which there is no LSP, BGP automatically uses a label-switched path along the path to reach the egress router.

To configure IS-IS so that it uses label-switched paths as shortcuts when installing information in the inet.3 routing table, include the shortcuts statement:

```
[edit protocols isis]
traffic-engineering {
  shortcuts;
}
```

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

Because the inet.3 routing table is present only on ingress routers, you can configure label-switched path shortcuts only on these routers.

For more information about configuring label-switched paths and MPLS, see the *JUNOS MPLS Applications Configuration Guide*.

Disabling IS-IS Support for Traffic Engineering

By default, IS-IS supports traffic engineering by exchanging basic information with the TED. To disable this support, and to disable IS-IS shortcuts if they are configured, include the `disable` statement:

```
[edit protocols isis]
traffic-engineering {
  disable;
}
```

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

Configuring the BFD Protocol

The bidirectional forwarding detection (BFD) protocol is a simple hello mechanism that detects failures in a network. 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. These timers are also adaptive and can be adjusted to be more or less aggressive.

To enable failure detection, include the `bfd-liveness-detection` statement:

```
[edit protocols isis interface interface-name]
bfd-liveness-statement {
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  minimum-transmit-interval milliseconds;
  multiplier number;
}
```

To specify the minimum transmit and receive interval for failure detection, include the `minimum-interval` statement:

```
[edit protocols isis interface interface-name]
minimum-interval milliseconds;
```

To specify the minimum receive interval for failure detection, include the `minimum-receive-interval` statement:

```
[edit protocols isis interface interface-name]
minimum-receive-interval milliseconds;
```

To specify the minimum transmit interval for failure detection, include the `minimum-transmit-interval` statement:

```
[edit protocols isis interface interface-name]
minimum-transmit-interval milliseconds;
```

To specify the detection time multiplier for failure detection, include the multiplier statement:

```
[edit protocols isis interface interface-name]
multiplier number;
```

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

Configuring Loose Authentication Check

To allow the use of MD5 authentication without requiring network-wide deployment, include the loose-authentication-check statement:

```
[edit protocols isis]
loose-authentication-check;
```

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

Disabling IS-IS

To disable IS-IS on the router without removing the IS-IS configuration statements from the configuration, include the disable statement:

```
[edit protocols]
isis {
  disable;
}
```

For a list of hierarchy levels at which you can configure 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 disable
[edit protocols]
user@host# show
isis;
```

Disabling IPv4 Routing

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

- Router does not advertise the NLPID for IPv4 in JUNOS software 0th LSP fragment

- Router does not advertise any IPv4 prefixes in JUNOS software LSPs

- Router does not advertise the NLPID for IPv4 in JUNOS software hello packets

Router does not advertise any IPv4 addresses in JUNOS software hello packets

Router does not calculate any IPv4 routes

To disable IPv4 routing on the router, include the `no-ipv4-routing` statement:

```
[edit protocols]
isis {
  no-ipv4-routing;
}
```

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

To re-enable IS-IS, remove the `no-ipv4-routing` statement from the configuration:

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

Disabling IPv6 Routing

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

Router does not advertise the NLPID for IPv6 in JUNOS software 0th LSP fragment

Router does not advertise any IPv6 prefixes in JUNOS software LSPs

Router does not advertise the NLPID for IPv6 in JUNOS software hello packets

Router does not advertise any IPv6 addresses in JUNOS software hello packets

Router does not calculate any IPv6 routes

To disable IPv6 routing on the router, include the `no-ipv6-routing` statement:

```
[edit protocols]
isis {
  no-ipv6-routing;
}
```

For a list of hierarchy levels at which you can configure 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
```

Configuring IS-IS Routing Policy

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 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:

```
[edit protocols isis]
export [ policy-names ];
```

For a list of hierarchy levels at which you can configure 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 Border Gateway Protocol (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 IS-IS Multicast Topologies

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.

To enable an alternate IPv4 multicast topology for IS-IS, include the `ipv4-multicast` statement:

```
[edit protocols]
isis {
  topologies {
    ipv4-multicast;
  }
}
```

To configure the multicast metric for an alternate multicast topology, include the `ipv4-multicast-metric` statement:

```
[edit protocols]
isis {
  interface interface-name {
    level level-number {
      ipv4-multicast-metric number;
    }
  }
}
```

To disable alternate multicast topologies for IS-IS, include the `no-ipv4-multicast` statement:

```
[edit protocols]
isis {
  interface interface-name {
    no-ipv4-multicast;
  }
}
```

To enable an alternate IPv6 multicast topology for IS-IS, include the `ipv6-multicast` statement:

```
[edit protocols]
isis {
  topologies {
    ipv6-multicast;
  }
}
```

To configure the multicast metric for an alternate IPv6 multicast topology, include the `ipv6-multicast-metric` statement:

```
[edit protocols]
isis {
  interface interface-name {
    level level-number {
      ipv6-multicast-metric number;
    }
  }
}
```

To disable alternate IPv6 multicast topologies for IS-IS, include the `no-ipv4-multicast` statement:

```
[edit protocols]
isis {
  interface interface-name {
    no-ipv6-multicast;
  }
}
```

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

Example: Configuring IS-IS Multicast Topologies

```
[edit]
protocols {
  isis {
    traceoptions {
      file isis size 5m world-readable;
    }
    isis so-0/0/0.0 {
      level 1 {
        metric 15;
        multicast-metric 18;
      }
      level 2 {
        metric 20;
        multicast-metric 14;
      }
    }
    isis so-1/0/0.0 {
      level 1 {
        metric 15;
        multicast-metric 12;
      }
      level 2 {
        metric 20;
        multicast-metric 23;
      }
    }
    isis so-2/0/0.0 {
      no-multicast;
      level 1 metric 14;
      level 2 metric 23;
    }
    isis fxp0.0 {
      disable;
    }
  }
}
```

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:

```
[edit protocols]
isis {
  topologies {
    ipv6-unicast;
  }
}
```

To configure a metric for an alternate IPv6 unicast topology, include the `ipv6-unicast-metric` statement:

```
[edit protocols]
isis {
  interface interface-name {
    level level-number {
      ipv6-unicast-metric number;
    }
  }
}
```

To disable alternate IPv6 unicast topologies for IS-IS, include the `no-ipv6-unicast` statement:


```
[edit protocols]
isis {
  interface interface-name {
    no-ipv6-unicast;
  }
}
```

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

Installing Default Route to Nearest Level 1/Level 2 Router

When a Level 1/Level 2 router (router B) finds out that it can reach at least one area other than its own (for example, in Area Y), it sets the ATTACHED bit in its Level 1 LSP. Thereafter, the Level 1 router (router A) introduces a default route pointing to the nearest attached Level 1/Level 2 router (router B). See Figure 5.

Figure 5: Install Default Route to Nearest Level 1/Level 2 Router



1461

Tracing IS-IS Protocol Traffic

To trace IS-IS protocol traffic, you can specify options in the global traceoptions statement at the [edit routing-options] hierarchy level, and you can specify IS-IS-specific options by including the traceoptions statement:

```
[edit protocols isis]
traceoptions {
  file name <replace> <size size> <files number> <no-stamp>
    <(world-readable | no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}
```

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

You can specify the following IS-IS-specific trace options in the IS-IS flag statement:

- all—Everything
- csn—Complete sequence number PDU (CSNP) packets
- error—Errored packets
- general—General events
- hello—Hello packets
- lsp—Link-state PDU (LSP) packets
- lsp-generation—Link-state PDU generation packets

normal—Normal events

packets—All IS-IS protocol packets

policy—Policy processing

psn—Partial sequence number PDU (PSNP) packets

route—Routing information

spf—Shortest-path-first (SPF) calculations

state—State transitions

task—Routing protocol task processing

timer—Routing protocol timer processing

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 traceoption flags detail and all with caution. These flags may cause the CPU to become very busy.

For information about tracing and global tracing options, see “Tracing Global Routing Protocol Operations” on page 99.

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;
    }
  }
}
```