

IS-IS

This chapter describes the Intermediate System to Intermediate System (IS-IS) routing protocol.

- About IS-IS, on page 1
- Prerequisites for IS-IS, on page 7
- Guidelines for IS-IS, on page 7
- Configure IS-IS, on page 8
- Monitoring IS-IS, on page 23
- History for IS-IS, on page 23

About IS-IS

IS-IS routing protocol is a link state Interior Gateway Protocol (IGP). Link-state protocols are characterized by the propagation of the information required to build a complete network connectivity map on each participating device. That map is then used to calculate the shortest path to destinations. The IS-IS implementation supports IPv4 and IPv6.

You can divide a routing domain into one or more subdomains. Each subdomain is called an area and is assigned an area address. Routing within an area is known as Level-1 routing. Routing between Level-1 areas is known as Level-2 routing. A router is referred to as an Intermediate System (IS). An IS can operate at Level 1, Level 2, or both. ISes that operate at Level 1 exchange routing information with other Level-1 ISes in the same area. ISes that operate at Level 2 exchange routing information with other Level-2 routers regardless of whether they are in the same Level-1 area. The set of Level-2 routers and the links that interconnect them form the Level-2 subdomain, which must not be partitioned in order for routing to work properly.

About NET

An IS is identified by an address known as a Network Entity Title (NET). The NET is the address of a Network Service Access Point (NSAP), which identifies an instance of the IS-IS routing protocol running on an IS. The NET is 8 to 20 octets in length and has the following three parts:

- Area address—This field is 1 to 13 octets in length and is composed of high-order octets of the address.

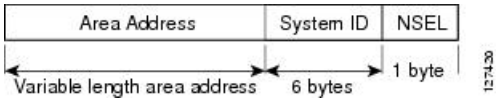
Note You can assign multiple area addresses to an IS-IS instance; in this case, all area addresses are considered synonymous. Multiple synonymous area addresses are useful when merging or splitting areas in the domain. Once the merge or split has been completed, you do not need to assign more than one area address to an IS-IS instance.

- System ID—This field is 6 octets long and immediately follows the area address. When the IS operates at Level 1, the system ID must be unique among all the Level-1 devices in the same area. When the IS operates at Level 2, the system ID must be unique among all devices in the domain.

Note You assign one system ID to an IS instance.

- NSEL—The N-selector field is 1 octet in length and immediately follows the system ID. It must be set to 00.

Figure 1: NET Format



IS-IS Dynamic Hostname

In the IS-IS routing domain, the system ID is used to represent each ASA. The system ID is part of the NET that is configured for each IS-IS ASA. For example, an ASA with a configured NET of 49.0001.0023.0003.000a.00 has a system ID of 0023.0003.000a. ASA-name-to-system-ID mapping is difficult for network administrators to remember during maintenance and troubleshooting on the ASAs.

The dynamic hostname mechanism uses link-state protocol (LSP) flooding to distribute the ASA-name-to-system-ID mapping information across the entire network. Every ASA on the network will try to install the system ID-to-ASA name mapping information in its routing table.

If an ASA that has been advertising the dynamic name type, length, value (TLV) on the network suddenly stops the advertisement, the mapping information last received will remain in the dynamic host mapping table for up to one hour, allowing the network administrator to display the entries in the mapping table during a time when the network experiences problems.

IS-IS PDU Types

ISes exchange routing information with their peers using protocol data units (PDUs). Intermediate System-to-Intermediate System Hello PDUs (IIHs), Link-State PDUs (LSPs), and Sequence Number PDUs (SNPs) types of PDUs are used.

IIHs

IIHs are exchanged between IS neighbors on circuits that have the IS-IS protocol enabled. IIHs include the system ID of the sender, the assigned area address(es), and the identity of neighbors on that circuit that are known to the sending IS. Additional optional information can also be included.

There are two types of IIHs:

- Level-1 LAN IIHs—These are sent on multiaccess circuits when the sending IS operates as a Level-1 device on that circuit.
- Level-2 LAN IIHs—These are sent on multiaccess circuits when the sending IS operates as a Level-2 device on that circuit.

LSPs

An IS generates LSPs to advertise its neighbors and the destinations that are directly connected to the IS. An LSP is uniquely identified by the following:

- System ID of the IS that generated the LSP
- Pseudonode ID—This value is always 0 except when the LSP is a pseudonode LSP
- LSP number (0 to 255)
- 32-bit sequence number

Whenever a new version of an LSP is generated, the sequence number is incremented.

Level-1 LSPs are generated by ISs that support Level 1. The Level-1 LSPs are flooded throughout the Level-1 area. The set of Level-1 LSPs generated by all Level-1 ISs in an area is the Level-1 LSP Database (LSPDB). All Level-1 ISs in an area have an identical Level-1 LSPDB and therefore have an identical network connectivity map for the area.

Level-2 LSPs are generated by ISs that support Level 2. Level-2 LSPs are flooded throughout the Level-2 subdomain. The set of Level-2 LSPs generated by all Level-2 ISs in the domain is the Level-2 LSP Database (LSPDB). All Level-2 ISs have an identical Level-2 LSPDB and therefore have an identical connectivity map for the Level-2 subdomain.

SNPs

SNPs contain a summary description of one or more LSPs. There are two types of SNPs for both Level 1 and Level 2:

- Complete Sequence Number PDUs (CSNPs) are used to send a summary of the LSPDB that an IS has for a given level.
- Partial Sequence Number PDUs (PSNPs) are used to send a summary of a subset of the LSPs for a given level that an IS either has in its database or needs to obtain.

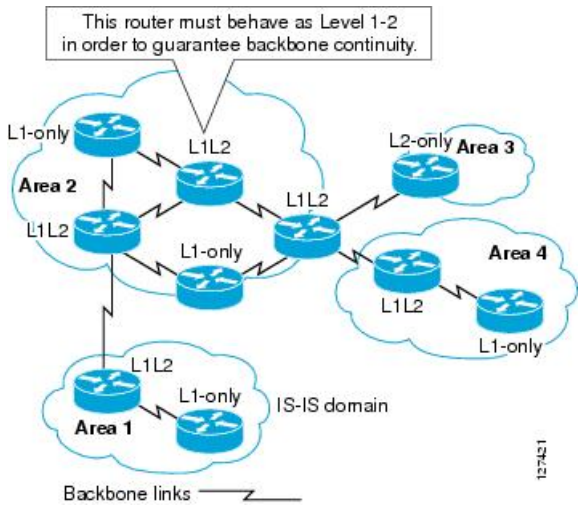
Operation of IS-IS on Multiaccess Circuits

Multiaccess circuits support multiple ISes, that is, two or more operating on the circuit. For multiaccess circuits a necessary prerequisite is the ability to address multiple systems using a multicast or broadcast address. An IS that supports Level 1 on a multiaccess circuit sends Level-1 LAN IIHs on the circuit. An IS that supports Level 2 on a multiaccess circuit sends Level-2 LAN IIHs on the circuit. ISes form separate adjacencies for each level with neighbor ISes on the circuit.

An IS forms a Level-1 adjacency with other ISes that support Level 1 on the circuit and has a matching area address. Two ISes with disjointed sets of area addresses supporting Level 1 on the same multiaccess circuit is NOT supported. An IS forms a Level-2 adjacency with other ISes that support Level 2 on the circuit.

The devices in the IS-IS network topology in the following figure perform Level 1, Level 2, or Level 1 and 2 routing along the backbone of the network.

Figure 2: Level-1, Level-2, Level 1-2 Devices in an IS-IS Network Topology



IS-IS Election of the Designated IS

If each IS advertised all of its adjacencies on a multiaccess circuit in its LSPs, the total number of advertisements required would be N 2 (where N is the number of ISes that operate at a given level on the circuit). To address this scalability issue, IS-IS defines a pseudonode to represent the multiaccess circuit. All ISes that operate on the circuit at a given level elect one of the ISes to act as the Designated Intermediate System (DIS) on that circuit. A DIS is elected for each level that is active on the circuit.

The DIS is responsible for issuing pseudonode LSPs. The pseudonode LSPs include neighbor advertisements for all of the ISes that operate on that circuit. All ISes that operate on the circuit (including the DIS) provide a neighbor advertisement to the pseudonode in their non-pseudonode LSPs and do not advertise any of their neighbors on the multiaccess circuit. In this way the total number of advertisements required varies as a function of N-the number of ISes that operate on the circuit.

A pseudonode LSP is uniquely classified by the following identifiers:

- System ID of the DIS that generated the LSP
- Pseudonode ID (ALWAYS NON-ZERO)
- LSP number (0 to 255)
- 32-bit sequence number

The nonzero pseudonode ID is what differentiates a pseudonode LSP from a non-pseudonode LSP and is chosen by the DIS to be unique among any other LAN circuits for which it is also the DIS at this level.

The DIS is also responsible for sending periodic CSNPs on the circuit. This provides a complete summary description of the current contents of the LSPDB on the DIS. Other ISes on the circuit can then perform the

following activities, which efficiently and reliably synchronizes the LSPDBs of all ISes on a multiaccess circuit:

- Flood LSPs that are absent from or are newer than those that are described in the CSNPs sent by the DIS.
- Request an LSP by sending a PSNP for LSPs that are described in the CSNPs sent by the DIS that are absent from the local database or older than what is described in the CSNP set.

IS-IS LSPDB Synchronization

Proper operation of IS-IS requires a reliable and efficient process to synchronize the LSPDBs on each IS. In IS-IS this process is called the update process. The update process operates independently at each supported level. Locally generated LSPs are always new LSPs. LSPs received from a neighbor on a circuit may be generated by some other IS or may be a copy of an LSP generated by the local IS. Received LSPs can be older, the same age, or newer than the current contents of the local LSPDB.

Handling Newer LSPs

When a newer LSP is added to the local LSPDB, it replaces an older copy of the same LSP in the LSPDB. The newer LSP is marked to be sent on all circuits on which the IS currently has an adjacency in the UP state at the level associated with the newer LSP—excluding the circuit on which the newer LSP was received.

For multiaccess circuits, the IS floods the newer LSP once. The IS examines the set of CNSPs that are sent periodically by the DIS for the multiaccess circuit. If the local LSPDB contains one or more LSPs that are newer than what is described in the CSNP set (this includes LSPs that are absent from the CSNP set), those LSPs are reflooded over the multiaccess circuit. If the local LSPDB contains one or more LSPs that are older than what is described in the CSNP set (this includes LSPs described in the CSNP set that are absent from the local LSPDB), a PSNP is sent on the multiaccess circuit with descriptions of the LSPs that require updating. The DIS for the multiaccess circuit responds by sending the requested LSPs.

Handling Older LSPs

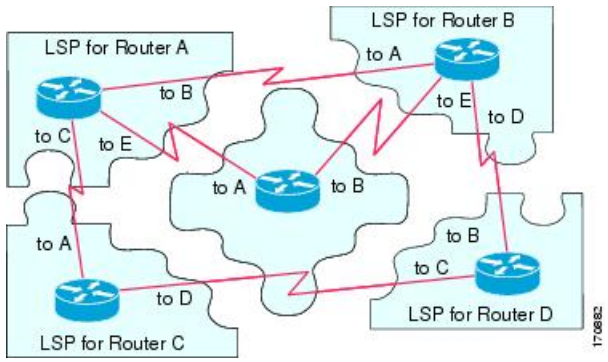
An IS may receive an LSP that is older than the copy in the local LSPDB. An IS may receive an SNP (complete or partial) that describes an LSP that is older than the copy in the local LSPDB. In both cases the IS marks the LSP in the local database to be flooded on the circuit on which the older LSP or SNP that contained the older LSP was received. Actions taken are the same as described above after a new LSP is added to the local database.

Handling Same-Age LSPs

Because of the distributed nature of the update process, it is possible than an IS may receive copies of an LSP that is the same as the current contents of the local LSPDB. In multiaccess circuits receipt of a same-age LSP is ignored. Periodic transmission of a CSNP set by the DIS for that circuit serves as an implicit acknowledgment to the sender that the LSP has been received.

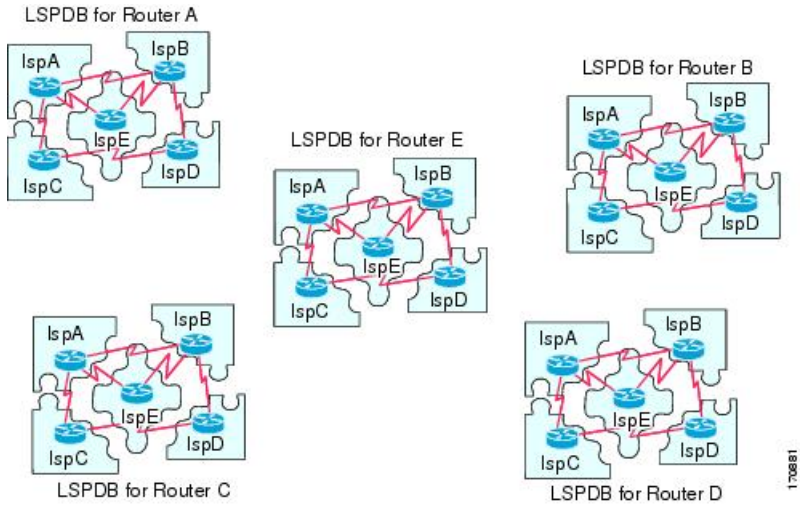
The following figure shows how LSPs are used to create a network map. Think of the network topology as a jigsaw puzzle. Each LSP (representing an IS) is one of the pieces. It is applicable to all Level-1 devices in an area or to all Level-2 devices in a Level-2 subdomain.

Figure 3: IS-IS Network Map



The following figure shows each device in the IS-IS network with its fully updated link-state database after the adjacencies have been formed among the neighbor devices. It is applicable to all Level-1 devices in an area or to all Level-2 devices in a Level-2 subdomain.

Figure 4: IS-IS Devices with Synchronized LSPDBs



IS-IS Shortest Path Calculation

When the contents of the LSPDB change, each IS independently reruns a shortest path calculation. The algorithm is based on the well-known Dijkstra algorithm for finding the shortest paths along a directed graph where the ISes are the vertices of the graph and the links between the ISes are edges with a nonnegative weight. A two-way connectivity check is performed before considering a link between two ISes as part of the graph. This prevents the use of stale information in the LSPDB, for example, when one IS is no longer operating in the network but did not purge the set of LSPs that it generated before stopping operation.

The output of the SPF is a set of tuples (destination, next hop). The destinations are protocol-specific. Multiple equal-cost paths are supported, in which case multiple next hops would be associated with the same destination.

Independent SPFs are performed for each level supported by the IS. When the same destination is reachable by both Level-1 and Level-2 paths, the Level-1 path is preferred.

A Level-2 IS that indicates that it has one or more Level-2 neighbors in other areas may be used by Level-1 devices in the same area as the path of last resort, also called the default route. The Level-2 IS indicates its attachment to other areas by setting an attached bit (ATT) in its Level-1 LSP 0.



Note An IS can generate up to 256 LSPs at each level. The LSPs are identified by the numbers 0 through 255. LSP 0 has special properties, including the significance of the setting of the ATT bit to indicate attachment to other areas. When LSPs that are numbered 1 through 255 have the ATT bit set, it is not significant.

IS-IS Shutdown Protocol

You can shut down IS-IS (placing it in an administrative down state) to make changes to the IS-IS protocol configuration without losing your configuration parameters. You can shut down IS-IS at the global IS-IS process level or at the interface level. If the device was rebooted when the protocol was turned off, the protocol would be expected to come back up in the disabled state. When the protocol is set to the administrative down state, network administrators are allowed to administratively turn off the operation of the IS-IS protocol without losing the protocol configuration, to make a series of changes to the protocol configuration without having the operation of the protocol transition through intermediate-and perhaps undesirable-states, and to then reenable the protocol at a suitable time.

Prerequisites for IS-IS

The following prerequisites are necessary before configuring IS-IS:

- Knowledge of IPv4 and IPv6.
- Knowledge of your network design and how you want traffic to flow through it before configuring IS-IS.
- Define areas, prepare an addressing plan for the devices (including defining the NETs), and determine the interfaces that will run IS-IS.
- Before you configure your devices, prepare a matrix of adjacencies that shows what neighbors should be expected in the adjacencies table. This will facilitate verification.

Guidelines for IS-IS

Firewall Mode Guidelines

Supported only in routed firewall mode. Transparent firewall mode is not supported.

Cluster Guidelines

Supported only in Individual Interface mode; Spanned EtherChannel mode is not supported.

Additional Guidelines

IS-IS is not supported with bidirectional forwarding.

Configure IS-IS

This section describes how to enable and configure the IS-IS process on your system.

Procedure

Step 1 Enable IS-IS Routing Globally, on page 8.
Step 2 Enable IS-IS Authentication, on page 9.
Step 3 Configure IS-IS LSP, on page 10
Step 4 Configure IS-IS Summary Addresses, on page 12.
Step 5 Configure IS-IS NET, on page 13.
Step 6 Configure IS-IS Passive Interfaces, on page 13.
Step 7 Configure IS-IS Interfaces, on page 14.
Step 8 Configure IS-IS IPv4 Address Family, on page 18.
Step 9 Configure IS-IS IPv6 Address Family, on page 21.

Enable IS-IS Routing Globally

Before you begin

In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **General**.
Step 2 Check the **Configure ISIS** check box to enable IS-IS.
Step 3 Check the **Shutdown protocol** check box to enable shutdown protocol,

See IS-IS Shutdown Protocol, on page 7for more information on Shutdown protocol.

Step 4 To have IS-IS use a dynamic hostname, check the **Use dynamic hostname** check box.

Dynamic hostname is enabled by default. See IS-IS Dynamic Hostname, on page 2 for detailed information on the dynamic hostname in IS-IS.

Step 5 To prevent IS-IS from padding LAN hello PDUs, check the **Do not pad LAN hello PDUs** check box.

IS-IS hellos are padded to the full maximum transmission unit (MTU) size. This allows for early detection of errors that result from transmission problems with large frames or errors that result from mismatched MTUs on adjacent interfaces. You can disable hello padding to avoid wasting network bandwidth in case the MTU of both interfaces is the same or in the case of translational bridging.

Step 6 To advertise passive interfaces only, check the **Advertise passive only** check box.

It excludes IP prefixes of connected networks from LSP advertisements, which reduces IS-IS convergence time.

Step 7 Choose whether to have your ASA act as station router (Level 1), an area router (Level 2), or both (Level 1-2) by clicking the appropriate radio button.

See About IS-IS, on page 1for more information on IS-IS levels.

Step 8 In the **Topology priority** field, enter a number that indicates where the ASA's priority is in the topology. The range is from 0 to 127.

Step 9 In the **Route priority tag** field, enter a tag that indicates the ASA's route priority. The range is from 1 to 4294967295. The default value is 100. Higher values indicate higher preference. This preference is sent to all routers in the IS-IS system.

Step 10 To have the IS conditionally advertise as L2, choose a device from the drop-down menu, and click **Manage**.

See Define a Route Map for the procedure for adding a route map.

Step 11 Check the **Log changes in adjacency** check box to have the ASA send a log message whenever an IS-IS neighbor goes up or down.

This command is disabled by default. Logging adjacency changes is useful when monitoring large networks.

Step 12 To have changes included from non-IIH events, check the **Include changes generated by non-IIH events** check box.

Step 13 To set up the skeptical time interval, enter the amount of minutes in the **Skeptical interval** field. The range is 0 to 1440 minutes. The default is five minutes.

Step 14 Click **Apply**.

Enable IS-IS Authentication

IS-IS route authentication prevents the introduction of unauthorized or false routing messages from unapproved sources. You can set a password for each IS-IS area or domain to prevent unauthorized routers from injecting false routing information into the link-state database, or you can configure a type of IS-IS authentication, either IS-IS MD5 or enhanced clear text authentication. You can also set authentication per interface. All IS-IS neighbors on interfaces configured for IS-IS message authentication must be configured with the same authentication mode and key for adjacencies to be established.

See About IS-IS, on page 1 for more information on areas and domains.

Before you begin

Before you can enable IS-IS route authentication, you must enable IS-IS and set up an area. See Enable IS-IS Routing Globally, on page 8 for the procedure.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **Authentication**.

Step 2 Configure authentication parameters for Level 1 and Level 2:

- In the **Key** field, enter the key to authenticate IS-IS updates. The key can include up to 16 characters.
- Click the **Enable** or **Disable** radio button depending or whether you want to have Send Only enabled.

  Note
  ASAs will have more time for the keys to be configured on each ASA if authentication is inserted only on the packets being sent, not checked on packets being received.

- Choose the authentication mode by clicking either the **Disabled**, **MD5**, or **Plaintext** radio button.

Step 3 If you choose **Disabled**, enter an area password for the Level 1 area (subdomain) and/or a domain password for the Level 2 domain.

Step 4 Click **Apply**.

Configure IS-IS LSP

An IS generates LSPs to advertise its neighbors and the destinations that are directly connected to IS-IS. See IS-IS PDU Types, on page 2 for more detailed information on LSPs.

Use the following commands to configure LSPs so that you have a faster convergence configuration.

Before you begin

In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **Link State Packet**.

Note
IS-IS must be enabled before you can configure LSP parameters. See Enable IS-IS Routing Globally, on page 8 for the procedure.

Step 2 To allow the ASA to ignore LSP packets that are received with internal checksum errors rather than purging the LSPs, check the **Ignore LSP errors** check box.

Step 3 To fast-flood and fill LSPs before running SPF, check the **Flood LSPs before running SPF**, and then in the **Number of LSPs to be flooded** field, enter a number. The range is 1 to 15. The default is 5.

This parameter sends a specified number of LSPs from the ASA. If no LSP number is specified, the default of 5 is used. The LSPs invoke SPF before running SPF. We recommend that you enable fast flooding, because then you speed up the LSP flooding process, which improves overall network convergence time.

Step 4 To suppress IP prefixes, check the **Suppress IP prefixes** check box, and then check one of the following:

- **Don't advertise IP prefixes learned form another ISIS level when ran out of LSP fragments**—Suppresses any routes coming from another level. For example, if the Level-2 LSP becomes full, routes from Level 1 are suppressed.

- **Don't advertise IP prefixes learned form other protocols when ran out of LSP fragments**—Suppresses any redistributed routes on the ASA.

In networks where there is no limit placed on the number of redistributed routes into IS-IS, it is possible that the LSP can become full and routes will be dropped. Use these options to control which routes are suppressed when the PDU becomes full.

Step 5 Configure the LSP generation intervals for Level 1 and Level 2:

- **LSP calculation interval**—Enter the interval of time in seconds between transmission of each LSP. The range is 1 to 120 seconds. The default is 5.

  The number should be greater than the expected round-trip delay between any two ASAs on the attached network. The number should be conservative or needless transmission results. Retransmissions occur only when LSPs are dropped. So setting the number to a higher value has little effect on reconvergence. The more neighbors the ASAs have, and the more paths over which LSPs can be flooded, the higher you can make this value.

- **Initial wait for LSP calculation**—Enter the time in milliseconds specifying the initial wait time before the first LSP is generated. The range is 1 to 120,000. The default is 50.

- **Minimum wait between first and second LSP calculation**—Enter the time in milliseconds between the first and second LSP generation. The range is 1 to 120,000. The default is 5000.

Step 6 If you want the values you configured for Level 1 to also apply to Level 2, check the **Use level 1 parameters also for level 2** check box.

Step 7 In the **Maximum LSP size** field, enter the maximum number of seconds between two consecutive occurrences of an LSP being generated. The range is 128 to 4352. The default is 1492.

Step 8 In the **LSP refresh interval** field, enter the number of seconds at which LSPs are refreshed. The range is 1 to 65,5535. The default is 900.

The refresh interval determines the rate at which the software periodically transmits in LSPs the route topology information that it originates. This is done to keep the database information from becoming too old.

Reducing the refresh interval reduces the amount of time that undetected link state database corruption can persist at the cost of increased link utilization. (This is an extremely unlikely event, however, because there are other safeguards against corruption.) Increasing the interval reduces the link utilization caused by the flooding of refreshed packets (although this utilization is very small).

Step 9 In the **Maximum LSP lifetime** field, enter the maximum number of seconds that LSPs can remain in a router's database without being refreshed. The range is 1 to 65,535. The default is 1200 (20 minutes).

You might need to adjust this parameter if you change the LSP refresh interval. LSPs must be periodically refreshed before their lifetimes expire. The value set for LSP refresh interval should be less that the value set for the maximum LSP lifetime; otherwise LSPs will time out before they are refreshed. If you make the LSP lifetime too low compared to the LSP refresh interval, the LSP refresh interval is automatically reduced to prevent the LSPs from timing out.

Step 10 Click **Apply**.

Configure IS-IS Summary Addresses

Multiple groups of addresses can be summarized for a given level. Routes learned from other routing protocols can also be summarized. The metric used to advertise the summary is the smallest metric of all the more specific routes. This helps to reduce the size of the routing table.

You need to manually define summary addresses if you want to create summary addresses that do not occur at a network number boundary or if you want to use summary addresses on an ASA with automatic route summarization disabled.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **Summary Address**.

The **Configure ISIS Summary Address** pane displays a table of the statically-defined IS-IS summary addresses. By default, IS-IS summarizes subnet routes to the network level. You can create statically defined IS-IS summary addresses to the subnet level from the **Configure ISIS Summary Address** pane.

Step 2 Click **Add** to add a new IS-IS summary address, or to click **Edit** to edit an existing IS-IS summary address in the table.

The **Add Summary Address** or **Edit Summary Address** dialog box is displayed. You can also double-click an entry in the table to edit that entry.

Step 3 In the **IP Address** field, enter the IP address of the summary route.

Step 4 In the **Netmask** field, choose or enter the network mask to apply to the IP address.

Step 5 Select the **Level 1**, **Level 2**, or **Level 1 and 2** radio button depending on which levels you want to receive summary addresses.

- **Level 1**—Summary routes are applied when redistributing routes into Level 1 and Level 2 and when Level 2 IS-IS advertises Level 1 routes as reachable in its area.
- **Level 2**—Routes learned by Level 1 routing are summarized into the Level 2 backbone with the configured address and mask value. Redistributed routes into Level 2 IS-IS are summarized also.
- **Level 1 and 2**—Summary routes are applied when redistributing routes into Level 1 and Level 2 and when Level 2 IS-IS advertises Level 1 routes as reachable in its area.

Step 6 In the **Tag** field, enter a number for the tag. The range is 1 to 4294967295.

The Tag field lets you specify a number to tag routes that are being summarized. If the routes have already been tagged on the **Configuration** > **Device Setup** > **Routing** > **ISIS** > **General** pane in the **Route priority tag** field, those routes are summarized, otherwise the tag is lost.

Step 7 In the **Metric** field, enter the metric that will be applied to the summary route. The range is 1 to 4294967295. The default value is 10.

The Metric value is assigned to the link and used to calculate the path cost via the links to destinations. You can configure this metric for Level 1 or Level 2 routing only.

Step 8 Click **OK**.

Step 9 Click **Apply**.

Configure IS-IS NET

Before you begin

In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

IS-IS uses addresses called Network Entity Titles (NET). They can be 8 to 20 bytes long, but are usually 10 bites long. You can add a NET entry on the NET page when clustering is not configured on the ASA. If your ASA has clustering configured, you must create a net pool entry on the **Configuration > Device Management > Advanced > Address Pools > NET Address Pools** pane. You can then reference the NET address pool on the NET pane.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **Network Entity Title (NET)**.

The **Configure Network Entity (NET)** pane displays a table of the NET addresses. You can add a NET entry here when clustering is NOT configured on the ASA. For an ASA with clustering configured, you must create a net pool entry at **Configuration** > **Device Management** > **Advanced** > **Address Pools** > **Net Address Pools**.

YOu can then reference the NET address pool on the Network Entity Title (NET) pane.

Step 2 Click **Add** to add a new IS-IS NET address, or to click **Edit** to edit an existing IS-IS NET address in the table.

The **Add Network Entity Title (NET)** or **Edit Network Entity Title (NET)** dialog box appears. You can also double-click an entry in the table to edit that entry.

Step 3 From the Network Entity Title (NET) drop-down list, choose a NET.

Step 4 In the **Maximum allowed Net** field, enter the maximum allowed NETs you want. The range is from 3 to 254. The default is 3.

In most cases only one NET is necessary, but in the case of merging multiple areas or splitting one area into multiple areas, you may need to use multiple area-addresses.

Step 5 Click **Apply**.

Configure IS-IS Passive Interfaces

You can disable IS-IS hello packets and routing updates on interfaces while still including the interface addresses in the topology database. These interfaces will not form IS-IS neighbor adjacencies

If you have an interface that you do not want to participate in IS-IS routing, but that is attached to a network that you want advertised, configure the passive interfaces to prevent that interface from using IS-IS. Additionally, you can specify the version of IS-IS that is used by the ASA for updates. Passive routing assists

in controlling the advertisement of IS-IS routing information and disables the sending and receiving of IS-IS routing updates on an interface.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **IS-IS** > **Passive Interfaces**.

Step 2 To suppress routing updates on all interfaces, check the **Suppress routing updates on all Interfaces** check box.

This causes all interfaces to operate in passive mode.

Step 3 To configure individual interfaces to suppress routing updates, select the named routing interface in the left column and click **Add** to add it to the Suppress routing updates column.

Specifying an interface name sets only that interface to passive mode. In passive mode, IS-IS routing updates are accepted by, but not sent out of, the specified interface.

Note
Only interfaces that you have given a dynamic hostname can be suppressed from sending routing updates. See IS-IS Dynamic Hostname, on page 2 for more information.

Step 4 Click **Apply**.

Configure IS-IS Interfaces

This procedure describes how to modify individual ASA interfaces for IS-IS routing.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **Interface**.

The **ISIS Interface Configuration** pane appears and displays the IS-IS interface configurations. You can configure hello padding per interface by checking/unchecking the **Hello Padding** check box.

IS-IS hellos are padded to the full maximum transmission unit (MTU) size. Padding IS-IS hellos to the full MTU allows for early detection of errors that result from transmission problems with large frames or errors that result from mismatched MTUs on adjacent interfaces.

Step 2 Choose an interface entry by double-clicking an interface entry, or choose the entry and click **Edit**.

The **Edit ISIS Interface** dialog box appears.

Step 3 On the **General** tab, configure the following:

- **Shutdown ISIS on this interface**—Lets you disable the IS-IS protocol for this interface without removing the configuration parameters. The IS-IS protocol does not form any adjacencies on this interface, and the IP address of this interface is put into the LSP that is generated by the ASA.
- **Enable ISIS on this interface**—Enables IS-IS protocol on this interface.

- **Enable IPv6 ISIS routing on this interface**—Enables IPv6 IS-IS routing on this interface.
- **Priority for level-1—**Lets you set a priority for Level 1. The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority becomes the DIS. The range is 0 to 127. The default is 64.

  Note
  In IS-IS, there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a router with a higher priority comes on line, it takes over the role from the current DIS. In the case of equal priorities, the highest MAC address breaks the tie.

- **Priority for level-2—**Lets you set a priority for Level 2. The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority becomes the DIS. The range is 0 to 127. The default is 64.

  Note
  In IS-IS, there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a router with a higher priority comes on line, it will take over the role from the current DIS. In the case of equal priorities, the highest MAC address breaks the tie.

- **Tag**—Sets a tag on the IP address configured for an interface when this IP prefix is put into an IS-IS LSP.
- **CSNP Interval for level-1—**Sets the Complete Sequence Number PDUs (CSNPs) interval in seconds between transmission of CSNPs on multiaccess networks for Level 1. This interval only applies for the designated ASA. The range is from 0 to 65535. The default is 10 seconds. It is unlikely that you will have to change the default.

  This option applies only for the designated router (DR) for a specified interface. Only DRs send CSNP packets to maintain database synchronization.

- **CSNP Interval for level-2—**Sets the Complete Sequence Number PDUs (CSNPs) interval in seconds between transmission of CSNPs on multiaccess networks for Level 2. This interval only applies for the designated ASA. The range is from 0 to 65535. The default is 10 seconds. It is unlikely that you will have to change the default.

  This option applies only for the designated router (DR) for a specified interface. Only DRs send CSNP packets to maintain database synchronization.

- **Adjacency filter—Filters the establishment of IS-IS adjacencies.**

  Filtering is performed by building NSAP addresses out of incoming IS-IS hello packets by combining each area address in the hello with the system ID. Each of these NSAP addresses is then passed through the filter. If any one NSAP matches, the filter is considered passed unless you specified**Match all area addresses**, in which case all addresses must pass. The functionality of **Match all area addresses** is useful in performing negative tests, such as accepting an adjacency only if a particular address is NOT present.

- **Match all area addresses—**(Optional) All NSAP addresses must match the filter to accept the adjacency. If not specified (the default), only one address must match the filter for the adjacency to be accepted.

Step 4 Click **OK**.

Step 5 On the **Authentication** tab, configure the following for Level 1 and/or Level 2:

- In the **Key** field, enter the key to authenticate IS-IS updates. The range is 0 to 8 characters.

  If no password is configured with the **Key** option, no key authentication is performed.

- For **Send only** click the **Enable** or **Disable** radio button.

  Choosing **Send only** causes the system only to insert the password into the SNPs, but not check the password in SNPs that it receives. Use this keyword during a software upgrade to ease the transition. The default is disabled.

- Choose the authentication mode by checking the **Mode** check box and then choosing **MD5** or **Text** from the drop-down list, and in the **Password** field, enter a password.

Step 6 Click **OK**.

Step 7 On the **Hello Padding** tab, configure the following:

- **Hello Padding**—Enables Hello Padding.

  IS-IS hellos are padded to the full maximum transmission unit (MTU) size. Padding IS-IS hellos to the full MTU allows for early detection of errors that result from transmission problems with large frames or errors that result from mismatched MTUs on adjacent interfaces.

- **Minimal holdtime 1 second for Level-1**—Enables the holdtime (in seconds) that the LSP remains valid for Level 1.

- **Hello Interval for level-1**—Specifies the length of time in seconds between hello packets for Level 1.

  By default, a value three times the hello interval seconds is advertised as the hold time in the hello packets sent. (Change the multiplier of 3 by checking the **Hello Multiplier** check box.) With smaller hello intervals, topological changes are detected faster, but there is more routing traffic. The range is 1 to 65535. The default is 10.

- **Minimal holdtime 1 second for Level-2**—Enables the holdtime (in seconds) that the LSP remains valid for Level 2.

- **Hello Interval for level-2**—Specifies the length of time in seconds between hello packets for Level 2.

  By default, a value three times the hello interval seconds is advertised as the hold time in the hello packets sent. (Change the multiplier of 3 by checking the **Hello Multiplier** check box.) With smaller hello intervals, topological changes are detected faster, but there is more routing traffic. The range is 1 to 65535. The default is 10.

- **Hello Multiplier for level-1**—Specifies the number of IS-IS hello packets a neighbor must miss before the ASA declares the adjacency is down for Level 1.

  The advertised hold time in IS-IS hello packets is set to the hello multiplier times the hello interval. Neighbors declare an adjacency to this ASA down after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can be set on a per-interface basis, and can be different between different ASAs in one area. The range is 3 to 1000. The default is 3.

- **Hello Multiplier for level-2**—Specifies the number of IS-IS hello packets a neighbor must miss before the ASA declares the adjacency is down for Level 2.

  The advertised hold time in IS-IS hello packets is set to the hello multiplier times the hello interval. Neighbors declare an adjacency to this ASA down after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can

be set on a per-interface basis, and can be different between different ASAs in one area. The range is 3 to 1000. The default is 3.

- **Configure Circuit Type**—Specifies whether the interface is configured for local routing (level 1), area routing (Level 2), or both local and area routing (Level 1-2).

Step 8 Click **OK**.

Step 9 On the **LSP Settings** tab, configure the following:

- **Advertise ISIS Prefix**—Allows the advertising of IP prefixes of connected networks in the LSP advertisements per IS-IS interface.

  Disabling this option is an IS-IS mechanism to exclude IP prefixed of connected network from LSP advertisements thereby reducing IS-IS convergence time.

- **Retransmit Interval**—Specifies the amount of time in seconds between retransmission of each IS-IS LSP.

  The number should be greater than the expected round-trip delay between any two ASAs on the attached network. The range is 0 to 65535. The default is 5.

- **Retransmit Throttle Interval**—Specifies the amount of time in milliseconds between retransmissions on each IS-IS LSP.

  This option may be useful in very large networks with many LSPs and many interfaces as a way of controlling LSP retransmission traffic. This option controls the rate at which LSPs can be resent on the interface. The range is 0 to 65535. The default is 33.

- **LSP Interval**—Specifies the time delay in millisecond between successive IS-IS LSP transmissions.

  In topologies with a large number of IS-IS neighbors and interfaces, an ASA may have difficulty with the CPU load imposed by LSP transmission and reception. This option allows the LSP transmission rate (and by implication the reception rate of other systems) to be reduced. The range is 1 to 4294967295. The default is 33.

Step 10 Click **OK**.

Step 11 On the **Metrics** tab, configure the following options for Level 1 and Level 2:

You can check the **Use the level 1 values also for level 2** check box if you want the same metrics for both Levels.

- **Use maximum metric value**— Specifies the metric assigned to the link and used to calculate the cost from each other router via the links in the network to other destinations.

- **Default metric**—Enter the number for the metric.

  The range is 1 to 16777214. The default value is 10.

Step 12 Click **OK**.

Step 13 Click **Apply**.

Configure IS-IS IPv4 Address Family

Routers are allowed to redistribute external prefixes or routes that are learned form any other routing protocol, static configuration, or connected interface. The redistributed routes are allowed in either a Level 1 router or a Level 2 router.

You can set up adjacency, Shortest Path First (SPF), and you can define conditions for redistributing routes from another routing domain into ISIS (redistribution) for IPv4 addresses.

Before you begin

Before you can enable IS-IS route authentication, you must enable IS-IS and set up an area. See Enable IS-IS Routing Globally, on page 8 for the procedure.

Make sure that IPv4 is enabled on at least one interface before trying to add a neighbor, or ASDM returns an error message indicating that the configuration failed.

Procedure

Step 1 Choose **Configuration > Device Setup > Routing > ISIS > IPv4 Address Family > General**.

a) Check the **Perform adjacency check** check box for the router to check on nearby IS routers.

b) In the **Administrative Distance** field, enter a distance assigned to routes discovered by IS-IS protocol.

Administrative distance is a parameter used to compare routes among different routing protocols. In general, the higher the value, the lower the trust rating. And administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. The range is 1 to 255. The default is 1.

You can use the distance option to configure the administrative distances applied to IS-IS routes when they are inserted into the Routing Information Base (RIB), and influence the likelihood of these routes being preferred over routes to the same destination addresses discovered by other protocols.

c) In the **Maximum number of forward paths** field, enter the maximum number of IS routes that can be installed in a routing table. The range is 1 to 8.

d) Check the **Distribute default route** check box to configure an IS routing process to distribute a default route, and then choose the default route from the drop-down list or click **Manage** to create a new route. See Define a Route Map for the procedure for creating a new route.

Step 2 Configure IS-IS metrics:

a) In the **Global ISIS metric for level 1**, enter a number specifying the metric.

The range is 1 to 63. The default is 10.

When you need to change the default metric value for all IS-IS interfaces, we recommend that you use the **Global ISIS metric for level 1** option to configure all interfaces globally. Globally configuring the metric values prevents user errors, such as unintentionally removing a set metric from an interface without configuring a new value and unintentionally allowing the interface to revert to the default metric of 10, thereby becoming a highly preferred interface in the network.

b) In the **Global ISIS metric for level 2**, enter a number specifying the metric.

The range is 1 to 63. The default is 10.

When you need to change the default metric value for all IS-IS interfaces, we recommend that you use the **Global ISIS metric for level 1** option to configure all interfaces globally. Globally configuring the metric values prevents user errors, such as unintentionally removing a set metric from an interface without configuring a new value and unintentionally allowing the interface to revert to the default metric of 10, thereby becoming a highly preferred interface in the network.

c) Choose one of the following to configure Type, Length, and Values (TLVs):

- Check the **Send and accept both styles of TLVs during transition** check box.
- Click the **Use old style of TLVs with narrow metric** radio button.
- Click the **Use new style TLVs to carry wider metric** radio button.

  If you choose one of the radio buttons, you can also check the **Accept both styles of TLVs during transition** check box.

  We strongly recommend that you use the new-style TLV because TLVs that are used to advertise IPv4 information in LSPs are defined to use only extended metrics. The software provides support of a 24-bit metric field, the wide metric. Using the new metric style, link metrics now have a maximum value of 16777214 with a total path metric of 4261412864.

d) Check the **Apply metric style to** check box, and then check the **Level-1** and/or **Level-2** check box.

Step 3 Click **Apply**.

Step 4 Choose **Configuration > Device Setup > Routing > ISIS > IPv4 Address Family > SPF**.

a) Check the **Honour external metrics during SPF calculations** check box, to have the SPF calculations include external metrics.

b) Check the **Signal other routers not to use this router as an intermediate hop in their SPF calculations** check box if you want to exclude this device, and configure the following:

- Check the **Specify on-startup behavior** check box, and choose one of the following:
  - **Advertise ourself as overloaded until BGP has converged**
  - **Specify time to advertise ourself as overloaded after reboot**

    In the **Time to advertise ourself as overloaded** field, enter the seconds to wait until the router advertises that it is overloaded. The range is 5 to 86400 seconds.
- Check the **Don't advertise IP prefixes learned from other protocols when overload bit it set** check box to exclude IP prefixes.
- Check the **Don't advertise IP prefixes learned from another ISIS level when overload bit is set** check box to exclude IP prefixes.

c) Configure the partial route calculation (PRC) intervals:

- In the **PRC Interval** field, enter an amount of time for the router to wait between partial route calculations (PRCs). The range is 1 to 120 seconds. The default is 5 seconds.
- In the **Initial wait for PRC** field, enter the initial PRC calculation delay (in milliseconds) after a topology change. The range is 1 to 120.000 milliseconds. The default is 2000 milliseconds.
- In the **Minimum wait between first and second PRC** field, enter the amount of time in milliseconds that you want the router to wait between PRCs. The range is 1 to 120,000 milliseconds. The default is 5000 milliseconds.

d) Configure the intervals for SPF calculations for Level 1 and Level 2:

Note
Check the **Use level 1 values also for level 2** check box if you want both levels to have the same values.

- In the **SPF Calculation Interval** field, enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.
- In the **Initial wait for SPF calculation** field, enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120.000 milliseconds. The default is 5500 milliseconds.
- In the **Minimum wait between first and second SPF calculation** field, enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.

Step 5 Click **Apply**.

Step 6 Choose **Configuration > Device Setup > Routing > ISIS > IPv6 Address Family > Redistribution**.

The **Redistribution** pane displays a table of the redistribution routes.

Step 7 Click **Add** to add a new redistribution route, or to click **Edit** to edit an redistribution route in the table.

The **Add Redistribution** or **Edit Redistribution** dialog box appears. You can also double-click an entry in the table to edit that entry.

a) From the **Source Protocol** drop-down list, choose the protocol (BGP, Connected, EIGRP, OSPF, RIP, or Static) from which you want to redistribute routes into the ISIS domain.
b) From the **Process ID** drop-down list, choose a process ID for the source protocol.
c) From the **Route Level** drop-down list, choose Level-1, Level- 2, or Level 1-2.
d) (Optional) In the **Metric** field, enter a metric for the redistributed route . The range is 1 to 4294967295.
e) For the **Metric Type**, click the internal or external radio button.
f) From the **Route Map** drop-down list, choose a route map that should be examined to filter the networks to be redistributed, or click **Manage** to add a new route map or edit an existing route map. See Define a Route Map for the procedure for configuring route maps.
g) Check one or more of the **Match** check boxes -Internal, External 1, External 2, NSSA External 1, and NSSA External 2 check boxes to redistribute routes from an OSPF network.

This step is only applicable for redistribution from OSPF networks.

Step 8 Click **OK**.

Step 9 Click **Apply**.

Attached Bit Configuration

In the following example, the attached-bit will stay set when the router matches 49.00aa in the L2 CLNS routing table:

```
ciscoasa(config)# router isis
 ciscoasa(config-router)# clns filter-set L2_backbone_connectivity permit 49.00aa
 ciscoasa(config-router)# route-map check-for-L2_backbone_connectivity
 ciscoasa(config-router)# match clns address L2_backbone_connectivity
ciscoasa(config)# router isis
 ciscoasa(config-router)#set-attached-bit route-map check-for-L2_backbone_connectivity
```

```
 ciscoasa(config-router)# end
 ciscoasa# show clns route 49.00aa

Known via "isis", distance 110, metric 30, Dynamic Entry
Routing Descriptor Blocks:
 via tr2, Serial0
   isis, route metric is 30, route version is 58
```

Configure IS-IS IPv6 Address Family

You can set up adjacency, SPF, and you can define conditions for redistributing routes from another routing domain into IS-IS (redistribution) for IPv6 addresses.

Before you begin

Before you can enable IS-IS route authentication, you must enable IS-IS and set up an area. See Enable IS-IS Routing Globally, on page 8 for the procedure.

Make sure that IPv6 is enabled on at least one interface before trying to add a neighbor, or ASDM returns an error message indicating that the configuration failed.

Procedure

Step 1 Choose **Configuration > Device Setup > Routing > ISIS > IPv6 Address Family > General**.

a) Check the **Perform adjacency check** check box for the router to check on nearby IS routers.

b) In the **Administrative Distance** field, enter a distance for the route. The range is 1 to 255. The default is 1.

Administrative distance is a parameter used to compare routes among different routing protocols. In general, the higher the value, the lower the trust rating. And administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. The range is 1 to 255. The default is 1.

You can use the distance option to configure the administrative distances applied to IS-IS routes when they are inserted into the Routing Information Base (RIB), and influence the likelihood of these routes being preferred over routes to the same destination addresses discovered by other protocols.

c) In the **Maximum number of forward paths** field, enter the maximum number of IS routes that can be installed in a routing table. The range is 1 to 8.

d) Check the **Distribute default route** check box to configure an IS routing process to distribute a default route, and then choose the default route from the drop-down list or click **Manage** to create a new route. See Define a Route Mapfor the procedure for creating a new route.

Step 2 Click **Apply**.

Step 3 Choose **Configuration > Device Setup > Routing > ISIS > IPv6 Address Family > SPF**.

a) Check the **Signal other routers not to use this router as an intermediate hop in their SPF calculations** check box if you want to exclude this device, and configure the following:

- Check the **Specify on-startup behavior** check box, and choose one of the following:
  - **Advertise ourself as overloaded until BGP has converged**
  - **Specify time to advertise ourself as overloaded after reboot**

In the **Time to advertise ourself as overloaded** field, enter the seconds to wait until the router advertises that it is overloaded. The range is 5 to 86,400 seconds.

- Check the **Don't advertise IP prefixes learned from other protocols when overload bit it set** check box to exclude IP prefixes.
- Check the **Don't advertise IP prefixes learned from another ISIS level when overload bit is set** check box to exclude IP prefixes.

b) Configure the partial route calculation (PRC) intervals: .

- In the **PRC Interval** field, enter an amount of time for the router to wait between partial route calculations (PRCs). The range is 1 to 120 seconds. The default is 5 seconds.
- In the **Initial wait for PRC** field, enter the amount of time for the router to wait for a PRC. The range is 1 to 120.000 milliseconds. The default is 2000 milliseconds.
- In the **Minimum wait between first and second PRC** field, enter the amount of time in milliseconds that you want the router to wait between PRCs. The range is 1 to 120.000 milliseconds. The default is 5000 milliseconds.

c) Configure the intervals for SPF calculations for Level 1 and Level 2:

Note
Check the **Use level 1 values also for level 2** check box if you want both levels to have the same values.

- In the **SPF Calculation Interval** field, enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.
- In the **Initial wait for SPF calculation** field, enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120.000 milliseconds. The default is 5500 milliseconds.
- In the **Minimum wait between first and second SPF calculation** field, enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.

Step 4 Click **Apply**.

Step 5 Choose **Configuration > Device Setup > Routing > ISIS > IPv6 Address Family > Redistribution**.

The **Redistribution** pane displays a table of the redistribution routes.

Step 6 Click **Add** to add a new redistribution route, or to click **Edit** to edit an redistribution route in the table.

The **Add Redistribution** or **Edit Redistribution** dialog box appears. You can also double-click an entry in the table to edit that entry.

a) From the **Source Protocol** drop-down list, choose the protocol (BGP, Connected, EIGRP, OSPF, RIP, or Static) from which you want to redistribute routes into the ISIS domain.
b) From the **Process ID** drop-down list, choose a process ID for the source protocol.
c) From the **Route Level** drop-down list, choose Level-1, Level- 2, or Level 1-2.
d) (Optional) In the **Metric** field, enter a metric for the redistributed route . The range is 1 to 4294967295.
e) For the **Metric Type**, click the **internal** or **external** radio button to specify the type of metric for the destination routing protocol.

f) From the **Route Map** drop-down list, choose a route map that should be examined to filter the networks to be redistributed, or click **Manage** to add a new route map or edit an existing route map. See Define a Route Map for the procedure for configuring route maps.

g) Check one or more of the **Match** check boxes -Internal, External 1, External 2, NSSA External 1, and NSSA External 2 check boxes to redistribute routes from an OSPF network.

This step is only applicable for redistribution from OSPF networks.

Step 7 Click **OK**.

Step 8 Click **Apply**.

Monitoring IS-IS

You can use the following screens to monitor the IS-IS routing process.

- **Monitoring > Routing > ISIS Neighbors** This pane shows information about each IS-IS neighbor.

  Each row represents one IS-IS neighbor. For each neighbor, the list includes the system ID, type, interface, IP address, the state (active, idle and so on), the hold time, and the circuit ID.

- **Monitoring > Routing > ISIS Rib** This pane displays the local IS-IS Routing Information Base (RIB) table.

- **Monitoring > Routing > ISIS IPv6 Rib** This pane displays the local IPv6 IS-IS RIB table.

History for IS-IS

Table 1: Feature History for IS-IS

Feature Name	Platform Releases	Feature Information
IS-IS routing	9.6(1)	The ASA now supports the Intermediate System to Intermediate System (IS-IS) routing protocol. Support was added for routing data, performing authentication, and redistributing and monitoring routing information using the IS-IS routing protocol. We introduced the following screens: **Configuration > Device Setup > Routing > ISIS** **Monitoring > Routing > ISIS**