



IS-IS Overview and Basic Configuration

This module provides a technical overview of the Integrated Intermediate System-to-Intermediate System (IS-IS) routing protocol. IS-IS 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 protocol was developed in the late 1980s by Digital Equipment Corporation (DEC) and was standardized by the International Standards Organization (ISO) in ISO/IEC 10589. The current version of this standard is ISO/IEC 10589:2002.

ISO/IEC 10589 defines support for the ISO Connectionless Network Protocol (CLNP) as defined in ISO 8473. However, the protocol was designed to be extensible to other network protocols. RFC 1195 defined IS-IS support for IP, and additional IETF extensions have defined IS-IS support for IPv6. Integration of support for multiple network layer protocols has led to the term Integrated IS-IS. The Cisco IOS IS-IS implementation supports CLNP, IPv4, and IPv6. This module and its related modules use the term IS-IS to refer to the Integrated IS-IS that is implemented by Cisco IOS software.

- [Prerequisites for IS-IS Overview and Basic Configuration, on page 1](#)
- [Information About IS-IS Overview and Basic Configuration, on page 2](#)
- [How to Create Monitor and Make Changes to a Basic IS-IS Network, on page 9](#)
- [Configuration Examples for a Basic IS-IS Network, on page 14](#)
- [Where to Go Next, on page 17](#)
- [Additional References for IS-IS Overview and Basic Configuration, on page 17](#)
- [Feature Information for IS-IS Overview and Basic Configuration, on page 18](#)
- [Glossary, on page 19](#)

Prerequisites for IS-IS Overview and Basic Configuration

- This document assumes knowledge of CLNS, IPv4, and IPv6.
- The amount of knowledge required for each technology is dependent on your deployment. You should know 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 Integrated IS-IS.
- To facilitate verification, a matrix of adjacencies should be prepared before you configure your devices, showing what neighbors should be expected in the adjacencies table.

Information About IS-IS Overview and Basic Configuration

IS-IS Functional Overview

A routing domain may be divided into one or more subdomains. Each subdomain is referred to as an area and is assigned an area address. Routing within an area is referred to as Level-1 routing. Routing between Level-1 areas is referred to as Level-2 routing. A device in Open Systems Interconnection (OSI) terminology is referred to as an Intermediate System (IS). An IS may operate at Level 1, Level 2, or both. ISs that operate at Level 1 exchange routing information with other Level-1 ISs in the same area. ISs that operate at Level 2 exchange routing information with other Level-2 devices regardless of whether they are in the same Level-1 area. The set of Level-2 devices and the links that interconnect them form the Level-2 subdomain, which must not be partitioned in order for routing to work properly.

IS Address Assignment

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 may be 8 to 20 octets in length and consists of three parts:

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



Note An IS-IS instance may be assigned multiple area addresses. When this is the case, all area addresses are considered synonymous. Multiple synonymous area addresses are useful when merging or splitting areas in the domain. In normal operation, for example, once the merge or split has been completed, there is no 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 An IS instance is assigned exactly one system ID.

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

The figure below shows the format for the NET.

Figure 1: NET Format



IS-IS PDU Types

ISs exchange routing information with their peers using protocol data units (PDUs). The following types of PDUs are used:

IIHs

Intermediate System-to-Intermediate System Hello PDUs (IIHs) are exchanged between IS neighbors on circuits on which the IS-IS protocol is 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 may also be included.

There are three types of IIHs:

- Point-to-Point IIHs—These are sent on point-to-point circuits.
- 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 Link-State PDUs (LSPs) to advertise its neighbors and the destination 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 (see “Operation of IS-IS on Multiaccess Circuits” section).
- 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 will have an identical Level-1 LSPDB and will 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 will have an identical Level-2 LSPDB and will therefore have an identical connectivity map for the Level-2 subdomain.

SNPs

Sequence Number PDUs (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.

For more information about how SNPs are used, see the “IS-IS Supported Circuit Types” section.

IS-IS Supported Circuit Types

IS-IS supports two generic circuit types:

- Point-to-point circuits
- Multiaccess circuits

Operation of IS-IS on Point-to-Point Circuits

A point-to-point circuit has exactly two ISs on the circuit. An IS forms a single adjacency to the other IS on the point-to-point circuit. The adjacency type describes what level(s) are supported on that circuit.

If both ISs support Level 1 on that circuit and the ISs are configured with at least one matching address, the adjacency supports Level 1. Level-1 LSPs and SNPs will be sent on that circuit.

If both ISs support Level 2 on that circuit, the adjacency supports Level 2. Level-2 LSPs and SNPs will be sent on that circuit.

The adjacency then can be Level 1, Level 2, or Level 1 and 2.

ISs send point-to-point IIHs on point-to-point circuits. These IIHs allow each IS to discover the identity of the neighbor, the configured area address(es), and the supported levels.

When an adjacency is first established, each IS sends a set of CSNPs for each level that is supported on the circuit. A CSNP set describes the current contents of the LSPDB at that level. By comparing the contents of the set of received CSNPs with the contents of the local LSPDB, each IS can determine where the databases differ and initiate procedures to exchange the necessary LSPs so that the databases are efficiently and reliably synchronized.

PSNPs are sent to acknowledge the receipt of an updated LSP.

Operation of IS-IS on Multiaccess Circuits

Multiaccess circuits support multiple ISs; for example, two or more operating on the circuit. The ability to address multiple systems utilizing a multicast or broadcast address is assumed.

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.


ISs form separate adjacencies for each level with neighbor ISs on the circuit.

An IS will form a Level-1 adjacency with other ISs that support Level 1 on the circuit and will have a matching area address. It is a misconfiguration to have two ISs with disjoint sets of area addresses supporting Level 1 on the same multiaccess circuit.

An IS will form a Level-2 adjacency with other ISs that support Level 2 on the circuit.

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

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



IS-IS Election of the Designated Intermediate System

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 ISs 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 ISs that operate on the circuit at a given level elect one of the ISs 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 ISs that operate on that circuit. All ISs 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 ISs 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 nonpseudonode 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 ISs on the circuit can then perform the following activities:

- Flood LSPs that they have 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.

In this way, the LSPDBs of all ISs on a multiaccess circuit are efficiently and reliably synchronized.

IS-IS Overview of 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. This section provides a brief overview of the operation of the update process. The update process operates independently at each supported level.

LSPs may be locally generated, in which case they always are new LSPs. LSPs may also be received from a neighbor on a circuit, in which case they may be generated by some other IS or may be a copy of an LSP generated by the local IS. Received LSPs may be older, the same age, or newer than the current contents of the local LSPDB.

Handling of Newer LSPs

A newer LSP is added to the local LSPDB. If an older copy of the same LSP currently exists in the LSPDB, it is replaced. 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.

On point-to-point circuits, the newer LSP will be flooded periodically until the neighbor acknowledges its receipt by sending a PSNP or by sending an LSP that is the same or newer than the LSP being flooded.

On multiaccess circuits, the IS will flood the newer LSP once. The IS examines the set of CSNPs 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 of 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.

At this point, the actions taken are identical to the actions that are described in the “Handling of Newer LSPs” section after a new LSP has been added to the local database.

Handling LSPs That Are the Same

Because of the distributed nature of the update process, it is possible that an IS may receive copies of an LSP that is the same as the current contents of the local LSPDB.

On a point-to-point circuit, receipt of such an LSP is ignored. Periodic transmission of a CSNP set by the DIS for that circuit will serve as an implicit acknowledgement to the sender that the LSP has been received.


In a multiaccess circuit, receipt of such an LSP is ignored. Periodic transmission of a CSNP set by the DIS for that circuit will serve as an implicit acknowledgement to the sender that the LSP has been received.

The figure below shows how the LSPs are used to create a network map. Imagine the network topology as a jigsaw puzzle. Each LSP (representing an IS) is considered one of the jigsaw pieces.



Note The figure below 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 figure below 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.



Note The figure below 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 Overview of the 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 ISs are the vertices of the graph and the links between the ISs are edges with a nonnegative weight. A two-way connectivity check is performed before considering a link between two ISs 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 ceasing operation.

The output of the SPF is a set of tuples (destination, next hop). The destinations are protocol-specific; for example, they would be prefixes when the supported protocol is IP, NSAPs of end systems when the supported protocol is CLNP. Multiple equal-cost paths are supported, in which case multiple next hops would be associated with the same destination.

Independent SPF calculations are performed for each level supported by the IS. In cases in which 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.

How to Create Monitor and Make Changes to a Basic IS-IS Network

Enabling IS-IS as an IP Routing Protocol on the Device

SUMMARY STEPS

1. `enable`
2. `configure terminal`
3. `router isis [area-tag]`
4. `net network-entity-title`
5. `end`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<p><code>enable</code></p> <p>Example:</p> <pre>Device> enable</pre>	<p>Enables privileged EXEC mode.</p> <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	<p><code>configure terminal</code></p> <p>Example:</p> <pre>Device# configure terminal</pre>	<p>Enters global configuration mode.</p>
Step 3	<p><code>router isis [area-tag]</code></p> <p>Example:</p> <pre>Device(config)# router isis</pre>	<p>Assigns a tag to an IS-IS process. Enters router configuration mode.</p> <ul style="list-style-type: none"> • Configure tags to identify multiple IS-IS processes by giving a meaningful name for each routing process. If the tag is not specified, a null tag (0) is assumed and the process is referenced with a null tag. The tag name must be unique among all IP router processes for the device.
Step 4	<p><code>net network-entity-title</code></p> <p>Example:</p> <pre>Device(config-router)# net 49.0001.0000.0000.000b.00</pre>	<p>Configures the NET on the device.</p> <ul style="list-style-type: none"> • The NET identifies the device for IS-IS.
Step 5	<p><code>end</code></p> <p>Example:</p>	<p>Exits router configuration mode and returns to privileged EXEC mode.</p>

	Command or Action	Purpose
	Device(config-router)# end	

Enabling IS-IS as an IP Routing Protocol on the Interface

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number*
4. **ip address** *ip-address mask [secondary]*
5. **ip router isis** [*area-tag*]
6. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	interface <i>type number</i> Example:	Enters interface configuration mode.
Step 4	ip address <i>ip-address mask [secondary]</i> Example: Device(config-if)# ip address 172.16.1.27 255.255.255.0	Sets the primary IP address on the interface.
Step 5	ip router isis [<i>area-tag</i>] Example: Device(config-if)# ip router isis company1	Enables IS-IS on the interfaces that are to use IS-IS to distribute their IP information (and additionally that might be used to establish IS-IS adjacencies). <ul style="list-style-type: none"> • Use the <i>area-tag</i> argument to specify to which IS-IS process the device belongs. • If there is more than one IS-IS process on the device, repeat the ip router isis command for each interface, specifying an area tag for each interface to associate each interface with the specific process to which it belongs.

	Command or Action	Purpose
Step 6	end Example: <pre>Device(config-if)# end</pre>	Exits interface configuration mode and returns to privileged EXEC mode.

Monitoring IS-IS

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **isis display delimiter** [**return** *count* | **character** *count*]
4. **exit**
5. **show ip protocols**
6. **show clns** *area-tag* **is-neighbors** [*type number*] [**detail**]
7. **show clns interface** [*type number*]
8. **show clns** *area-tag* **neighbors** [*type number*] [**area**] [**detail**]
9. **show clns** *area-tag* **traffic**
10. **show ip route** [*ip-address* [*mask*]] [[**longer-prefixes**] | *protocol* [*process-id*] | **list** [*access-list-number* | *access-list-name*] | **static download**]]
11. **show isis** [*process-tag*] **database** [**level-1**] [**level-2**] [**I1**] [**I2**] [**detail**] [**lspid**]
12. **show isis database verbose**
13. **show isis lsp-log**
14. **show isis** [*area-tag*] [**ipv6** | *] **spf-log**
15. **show isis** [*process-tag*] [**ipv6** | *] **topology**
16. **show isis** [*area-tag*] **neighbors** [**detail**]

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>Device> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <pre>Device# configure terminal</pre>	Enters global configuration mode.
Step 3	isis display delimiter [return <i>count</i> character <i>count</i>] Example: <pre>Device(config)# isis display delimiter return 3</pre>	Makes output from multiarea displays easier to read by specifying the delimiter to use to separate displays of information.

	Command or Action	Purpose
Step 4	exit Example: <pre>Device(config)# exit</pre>	Returns to privileged EXEC mode.
Step 5	show ip protocols Example: <pre>Device# show ip protocols</pre>	Displays the parameters and current state of the active routing protocol process. <ul style="list-style-type: none"> You can use this command to learn what protocols are active, what interfaces they are active on, what networks they are routing for, and other parameters that relate to the routing protocols.
Step 6	show clns area-tag is-neighbors [type number] [detail] Example: <pre>Device# show clns is-neighbors detail</pre>	Displays IS-IS information for IS-IS device adjacencies.
Step 7	show clns interface [type number] Example: <pre>Device# show clns interface</pre>	List the CLNS-specific information about each interface.
Step 8	show clns area-tag neighbors [type number] [area] [detail] Example: <pre>Device# show clns area3 neighbors</pre>	Displays both ES and IS neighbors. <ul style="list-style-type: none"> The show clns neighbor command output verifies that the right adjacencies have established. A matrix of adjacencies should be prepared before you configure your devices, showing what neighbors should be expected in the adjacencies table, to facilitate verification.
Step 9	show clns area-tag traffic Example: <pre>Device# show clns area3 traffic</pre>	Displays traffic statistics. <p>To monitor IS-IS for stability once it has been deployed across your network, enter the show clns traffic command to check the following important statistics: high numbers of SPFs, checksum errors, and retransmissions. To troubleshoot IS-IS behavior, you can use the output from the show clns traffic command to check for the following indicators:</p> <ul style="list-style-type: none"> The number of link-state PDUs (LSPs) can help you determine the stability of the IS-IS network. The number of LSPs should never be zero. However, an LSP count that keeps increasing over a short time period indicates a network issue. LSP retransmissions should stay low. A later execution of the show clns traffic command that shows an increase in LSP retransmissions, as

	Command or Action	Purpose
		<p>compared to an earlier execution of the command, can indicate instability or traffic problems.</p> <ul style="list-style-type: none"> To check for partial route calculations (PRCs), enter the show cns traffic command. PRCs are flooded when a change that does not affect topology is reported through an LSP; typical examples include the addition or removal of a prefix or metric changes for external or passive interfaces. A PRC update queue that remains full or increases to the maximum value for long periods of time indicates network instability. LSP checksum errors indicate a problem. The update queue should not stay full and should not drop much.
Step 10	<p>show ip route [<i>ip-address</i> [<i>mask</i>]] [[longer-prefixes] <i>protocol</i> [<i>process-id</i>] list [<i>access-list-number</i> <i>access-list-name</i>] static download]]</p> <p>Example:</p> <pre>Device# show ip route 172.16.0.21</pre>	Displays the current state of the routing table.
Step 11	<p>show isis [<i>process-tag</i>] database [level-1] [level-2] [11] [12] [detail] [lspid]</p> <p>Example:</p> <pre>Device# show isis database detail</pre>	<p>Displays additional information about the IS-IS database.</p> <ul style="list-style-type: none"> Displays the link-state database for Level-1 and Level-2, the contents for each LSP, and the link-state protocol PDU identifier.
Step 12	<p>show isis database verbose</p> <p>Example:</p> <pre>Device# show isis database verbose</pre>	Displays additional information about the IS-IS database such as the sequence number, checksum, and holdtime for LSPs.
Step 13	<p>show isis lsp-log</p> <p>Example:</p> <pre>Device# show isis lsp-log</pre>	Displays a log of LSPs including time of occurrence, count, interface, and the event that triggered the LSP.
Step 14	<p>show isis [<i>area-tag</i>] [ipv6 *] spf-log</p> <p>Example:</p> <pre>Device# show isis spf-log</pre>	<p>Displays how often and why the device has run a full shortest path first (SPF) calculation.</p> <ul style="list-style-type: none"> If the device continues to run SPF without ceasing, there might be an issue regarding a change in the network (intra-area). The cause for the continued SPF calculations could be an interconnecting link that is transitioning up/down/up/down or a metric change. It is normal for the SPF calculation to run a few times

	Command or Action	Purpose
		when a network change occurs, but then it should cease.
Step 15	show isis [<i>process-tag</i>] [ipv6 *] topology Example: Device# show isis topology	Displays a list of all connected devices in all areas.
Step 16	show isis [<i>area-tag</i>] neighbors [detail] Example: Device# show isis neighbors detail	Displays IS-IS adjacency information. <ul style="list-style-type: none"> The show isis neighbor detail command output verifies that the right adjacencies have established. A matrix of adjacencies should be prepared before you configure your devices, showing what neighbors should be expected in the adjacencies table, to facilitate verification.

Example

When the **show isis neighbors** command is entered with the **detail** keyword, the output provides information about the IS-IS adjacencies that have formed.

```
Device1# show isis neighbors detail

System Id      Type Interface IP Address      State Holdtime Circuit Id
Device2        L2 Et1/0         10.1.1.0        UP    255      Circuit3.01
Area Address(es): 32
SNPA: aabb.cc00.2001
State Changed: 00:00:14
LAN Priority: 64
Format: Phase V
```

Troubleshooting Tips

You can use the following two system debugging commands to check your IS-IS IPv4 implementation.

- If adjacencies are not coming up properly, use the **debug isis adj-packets** command.
- To display a log of significant events during an IS-IS SPF calculation, use the **debug isis spf-events** command.

Configuration Examples for a Basic IS-IS Network

Example: Configuring a Basic IS-IS Network

The following example shows how to configure three devices to run IS-IS as an IP routing protocol.

Device A Configuration

```

router isis
 net 49.0001.0000.0000.000a.00
interface ethernet0/0
 ip address 10.1.1.1 255.255.255.0
 ip router isis
interface serial 2/0
 ip router isis
 ip address 192.168.1.2 255.255.255.0

```

Device B Configuration

```

router isis
 net 49.0001.0000.0000.000b.00
interface ethernet0/0
 ip router isis
 ip address 172.17.1.1 255.255.255.0
interface serial2/0
 ip router isis
 ip address 192.168.1.1 255.255.255.0
interface serial5/0
 ip router isis
 ip address 172.21.1.1 255.255.255.0

```

Device C Configuration

```

router isis
 net 49.0001.0000.0000.000c.00
interface ethernet2/0
 ip router isis
 ip address 172.21.1.2 255.255.255.0
interface serial5/0
 ip router isis
 ip address 172.22.1.1 255.255.255.0

```

The **show isis topology** command displays the following information about how the devices are connected within the IS-IS network:

```
DeviceB# show isis topology
```

```

IS-IS paths to level-1 routers
System Id      Metric  Next-Hop      Interface  SNPA
DeviceA        10     DeviceA       Se2/0      *HDLC*
DeviceB        --
DeviceC        10     DeviceC       Se5/0      *HDLC*
IS-IS paths to level-2 routers
System Id      Metric  Next-Hop      Interface  SNPA
DeviceA        10     DeviceA       Se2/0      *HDLC*
DeviceB        --
DeviceC        10     DeviceC       Se5/0      *HDLC*

```

The **show isis database** command displays following information for the Level 1 and Level 2 LSPs for each device in the IS-IS network.

```
DeviceB# show isis database
```

```

IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL

```

Example: Configuring a Basic IS-IS Network

```

DeviceA.00-00      0x00000005  0x1A1D      1063      0/0/0
DeviceB.00-00      * 0x00000006  0xD15B      1118      0/0/0
DeviceC.00-00      0x00000004  0x3196      1133      1/0/0
IS-IS Level-2 Link State Database:
LSPID              LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
DeviceA.00-00      0x00000008  0x0BF4      1136      0/0/0
DeviceB.00-00      * 0x00000008  0x1701      1137      0/0/0
DeviceC.00-00      0x00000004  0x3624      1133      0/0/0

```

The **show ip route** command displays information about the interfaces of each device, including their IP addresses and how they are connected to Device B:

```

DeviceB# show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set
 172.17.0.0/24 is subnetted, 1 subnets
C       172.17.1.0 is directly connected, Ethernet0/0
 172.16.0.0/24 is subnetted, 1 subnets
C       172.16.1.0 is directly connected, Serial4/0
 172.21.0.0/24 is subnetted, 1 subnets
C       172.21.1.0 is directly connected, Serial5/0
 172.22.0.0/24 is subnetted, 1 subnets
i L1    172.22.1.0 [115/20] via 172.21.1.2, Serial5/0
 10.0.0.0/24 is subnetted, 1 subnets
i L1    10.1.1.0 [115/20] via 192.168.1.2, Serial2/0
C       192.168.1.0/24 is directly connected, Serial2/0
C       192.168.3.0/24 is directly connected, Serial3/0

```

The **show isis spf-log** command displays logs of Level 1 and Level 2 LSPs including time of occurrence, duration, count, and the event that triggered the LSP.

```


DeviceC## show isis spf-log

  level 1 SPF log
  When   Duration  Nodes  Count  First trigger LSP  Triggers
00:01:30      0       3       7      DeviceB.00-00  PERIODIC NEWADJ NEWLSP TLVT
  level 2 SPF log
  When   Duration  Nodes  Count  First trigger LSP  Triggers
00:01:31      0       3       7      DeviceB.00-00  PERIODIC NEWADJ NEWLSP TLVT

```

The figure below illustrates the sample configuration.

Figure 5: IS-IS Routing



Where to Go Next

- To initially configure and enable IS-IS, see the “Configuring a Basic IS-IS Network” module.
- To customize IS-IS for your network design, see the “Customizing IS-IS for Your Network Design” module.
- To customize IS-IS for achieving fast convergence and scalability, see the following modules:
 - “Overview of IS-IS Fast Convergence”
 - “Setting Best Practice Parameters for IS-IS Fast Convergence”
 - “Reducing Failure Detection Times in IS-IS Networks”
 - “Reducing Link Failure and Topology Change Notification Times in IS-IS Networks”
 - “Reducing Alternate-Path Calculation Times in IS-IS Networks”
- To enhance IS-IS network security, see the "Enhancing Security in an IS-IS Network" module.

Additional References for IS-IS Overview and Basic Configuration

Related Documents

Related Topic	Document Title
IPv6 Routing: IS-IS Support for IPv6	"IPv6 Routing: IS-IS Support for IPv6 " module
IPv6 Routing: Route Redistribution	"IPv6 Routing: Route Redistribution" module
IPv6 Routing: IS-IS Support for IPv6	"IPv6 Routing: IS-IS Support for IPv6 " module

Standards

Standard	Title
ISO 8473	<i>CLNP, Connectionless Network Protocol</i>
ISO 9542	<i>ES-IS Routing Information Exchange Protocol</i>
ISO/IEC 10589	<i>IS-IS Protocol</i>

MIBs

MIB	MIBs Link
No new or modified MIBs are supported, and support for existing MIBs has not been modified.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 1195	<i>Use of OSI IS-IS for Routing in TCP/IP and Dual Environments</i> (http://www.ietf.org/rfc/rfc1195.txt)

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for IS-IS Overview and Basic Configuration

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Glossary

area —A physically connected portion of a routing domain in which all devices are assigned a common area address. Also known as the Level-1 subdomain. A routing domain may consist of multiple areas that are reachable by traversing the Level-2 subdomain.

area address —The high-order octets of the Network Entity Title (NET) assigned to an IS. All ISs in the same Level-1 area are assigned the same area address.

CLNP —ISO Connectionless Network Protocol as defined in ISO 8473.

DIS —Designated Intermediate System. An IS elected by all the ISs operating on a multiaccess circuit at a given level to represent the multiaccess circuit. The DIS sends pseudonode LSPs on behalf of the circuit advertising adjacencies to all the ISs operating on that circuit.

domain —The portion of a network on which the IS-IS protocol is configured to operate. The routing domain consists of all Level-1 areas and the Level-2 subdomain.

ES —end system. An ES is any nonrouting host or node.

Integrated IS-IS —Extended form of IS-IS that supports multiple network protocols. Extensions have been defined in IETF documents, especially RFC 1195.

IS —intermediate system. OSI term for a device.

IP —Internet Protocol Version 4, also known as IPv4.

IPv6 —Internet Protocol Version 6.

IS-IS —Intermediate System-to-Intermediate System. Routing protocol as defined in ISO/IEC 10589.

Level-1 router —An IS that supports Level-1 routing for its assigned area.

Level-2 router —An IS that supports Level-2 routing.

Level-2 subdomain —All Level-2 capable devices in a domain and the links that interconnect them. Level-1 areas are interconnected via the Level-2 subdomain. For routing in a domain to work properly, the Level-2 subdomain must not be partitioned.

NET —Network Entity Title. An address assigned to an instance of the IS-IS protocol. The NET includes an area address, a system ID, and an N-selector. When multiple NETs are assigned to an IS-IS instance, only the area address portion of the NET may differ.

NSEL —N-selector. The least significant octet of a Network Entity Title. It is always assigned the value 00.

system ID —The part of the NET that immediately follows the area address. The field is 6 octets long.

